CONF-9411140-Absts.

Inaging Sciences Workshop

Proceedings

James V. Candy
C.A.S.I.S. Director

Livermore, CA
November 15, 1994

November 15, 1994

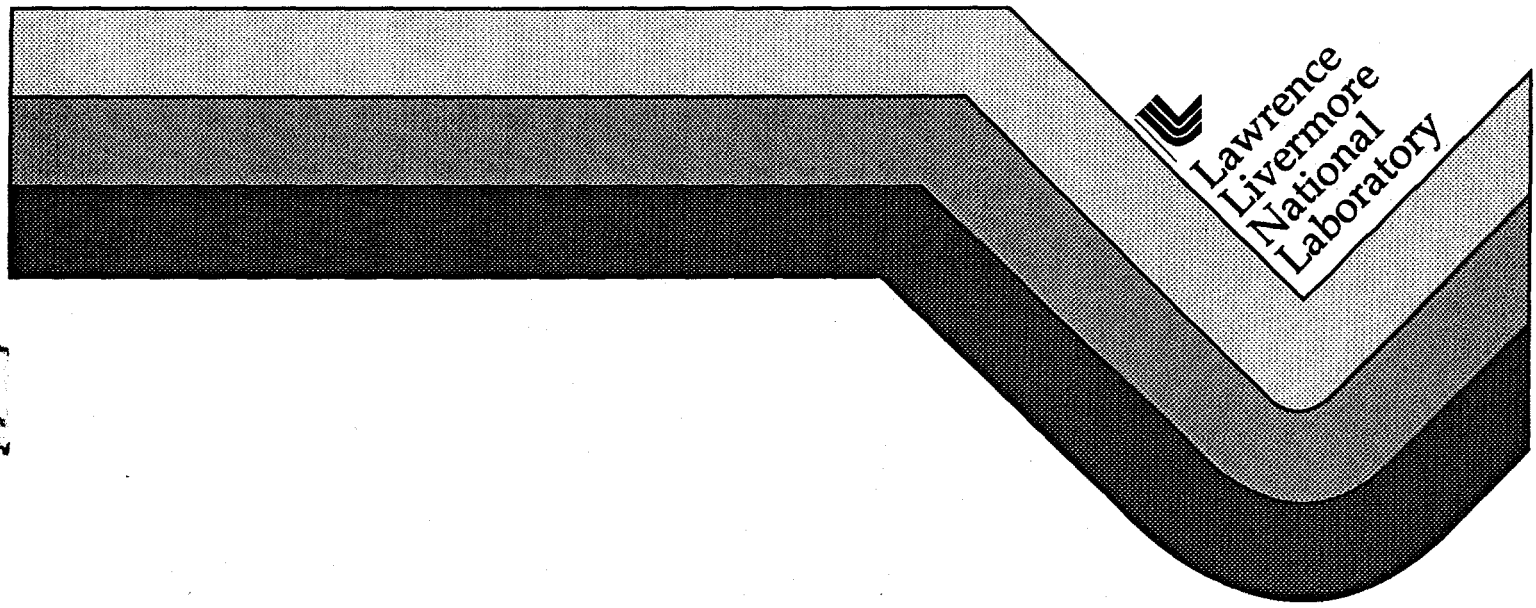

Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Imaging Sciences Workshop

Proceedings

James V. Candy

Director
Center for Advanced Signal and Image Sciences
November 15, 1994

Welcome

We would like to welcome you to the *Imaging Sciences Workshop* sponsored by ***C.A.S.I.S.***, the Center *for Advanced Signal & Image Sciences.* The Center was established primarily to provide a forum where researchers can freely exchange ideas and information in a comfortable intellectual environment focused on the areas of the signal and image sciences. This Workshop represents that forum for Lawrence Livermore National Laboratory personnel and invited speakers/participants to present problems and possible solutions to stimulate the exchange of ideas, papers, software, etc.

As you can see from the Technical Program, applications in the Signal & Image Sciences at LLNL are far reaching, ranging from astronomical imaging to microstructure imaging. Our program is organized, first, by imaging application areas: astronomical, biomedical, vision, acoustic, microwave, physical, tomography and then by function: algorithms, hardware and software.

We hope that you will enjoy the mix of presentations, meet the participating scientists and engineers and exchange your ideas openly. Thanks for attending.

J. V. Candy
C.A.S.I.S. Director

Table of Contents

Acoustic/Oceanic Imaging

Microwave/Acoustic Imaging

Computed Tomography

Physical Imaging

Imaging Algorithms

ASTRONOMICAL IMAGING

IMAGING SCIENCES WORKSHOP • 1994

Adaptive Optics and Laser Guide Stars at Lick Observatory

James M. Brase

Adaptive Optics at LLNL

For the past several years LLNL has been developing adaptive optics systems for correction of both atmospheric turbulence effects and thermal distortions in optics for high-power lasers. Our early work focused on adaptive optics for beam control in laser isotope separation and ground-based free electron lasers. We are currently developing innovative adaptive optics and laser systems for sodium laser guide star applications at the University of California's Lick and Keck Observeratories. This talk will describe our adaptive optics technology and some of its applications in high-resolution imaging and beam control.

Laser Guide Star

This talk will describe a project in which an artificial "star" is produced by shining a laser straight up into the air and exciting a layer of sodium atoms at an altitude of about 60 miles. A special telescope with "adaptive" optics then observes the sodium layer "star", corrects the image to account for atmospheric turbulence (i.e. the "twinkling" we normally see in stars), and then applies to use modest telescopes to achieve very high quality observational astronomy images.

Optical Design of the Comet Shoemaker-Levy Speckle Camera

Horst Bissinger

An optical design is presented in which the Lick 3 meter telescope and a bare CCD speckle camera system was used to image the collision sites of the Shoemaker -Levy 9 comet with the Planet Jupiter. The brief overview includes of the optical constraints and system layout. The choice of a Risley prism combination to compensate for the time dependent atmospheric chromatic changes are described. Plate scale and signal-to-noise ratio curves resulting from imaging reference stars are compared with theory. Comparisons between un-corrected and reconstructed images of Jupiter's impact sites. The results confirm that speckle imaging techniques can be used over an extended time period to provide a method to image large extended objects.

Speckle Imaging Algorithms for Planetary Imaging

Erik Johansson

I will discuss the speckle imaging algorithms used to process images of the impact sites of the collision of comet Shoemaker - Levy 9 with Jupiter. The algorithms use a phase retrieval process based on the average bispectrum of the speckle image data. High resolution images are produced by estimating the Fourier magnitude and Fourier phase of the image separately, then combining them and inverse transforming to achieve the final result. I will show raw speckle image data and high-resolution image reconstructions from our recent experiment at Lick Observatory.

Comet Shoemaker-Levy 9/Jupiter Collision Observed with a High Resolution Speckle Imaging System

Don Gavel

During the week of July 16, 1994, comet Shoemaker-Levy 9, broken into 20 plus pieces by tidal forces on its last orbit, smashed into the planet Jupiter, releasing the explosive energy of 500 thousand megatons.

A team of observers from LLNL used the LLNL Speckle Imaging Camera (see Horst Bissinger's talk) mounted on the University of California's Lick Observatory 3 Meter Telescope to capture continuous sequences of planet images during the comet encounter. Post processing with the bispectral phase reconstruction algorithm (described in the talk by Erik Johansson) improves the resolution by removing much of the blurring due to atmospheric turbulence. High resolution images of the planet surface showing the aftermath of the impacts are probably the best that were obtained from any ground-based telescope.

We have been looking at the regions of the fragment impacts to try to discern any dynamic behavior of the spots left on Jupiter's cloud tops. Such information can lead to conclusions about the nature of the comet and of Jupiter's atmosphere. So far, the Hubble Space Telescope has observed expanding waves from the G impact whose mechanism is enigmatic since they appear to be too slow to be sound waves and too fast to be gravity waves, given the present knowledge of Jupiter's atmosphere. Some of our data on the G and L impact region complements the Hubble observations but, so far, is inconclusive about spot dynamics.

In this talk I will present our latest images, including a movie of one night's observations, and will discuss the latest theories on the comet collision.

Optical Design of the Adaptive Optics Laser Guide Star System

Horst Bissinger

The design of an adaptive optics package for the 3 meter Lick telescope is presented. This instrument package includes a 69 actuator deformable mirror and a Hartmann type wavefront sensor operating in the visible wavelength; a quadrant detector for the tip-tilt sensor and a tip-tilt mirror to stabilize atmospheric first order tip-tilt errors. A high speed computer drives the deformable mirror to achieve near diffraction limited imagery. The different optical components and their individual design constraints are described. Motorized stages and diagnostics tools are used to operate and maintain alignment throughout observation time from a remote control room. The expected performance are summarized and actual results of astronomical sources are presented

Stellar Compass for the Clementine Mission

Bruce Wilson

A CCD sensor with 42 x 28 degrees FOV and 576 x 384 pixels was built by the Advanced Technology Program (ATP) in the Physics Department at LLNL. That sensor, called the StarTracker camera, is used on the Clementine Lunar Mapping mission between January and May, 1994. Together with the Stellar Compass software, the StarTracker camera provided a way of identifying its orientation to within about 150 microradians in camera body pitch and yaw.

This presentation will be an overview of basically how the Stellar Compass software works, along with showing some of its performance results.

BIOMEDICAL IMAGING

IMAGING SCIENCES WORKSHOP • 1994

Infrared Tomography for Diagnostic Imaging of Port Wine Stain Blood Vessels

Dennis Goodman

The objective of this work is the development of Infrared Tomography (IRT) for detecting and characterizing subsurface chromophores in human skin. Characterization of cutaneous chromophores is crucial for advances in the laser treatment of pigmented lesions (e.g., port wine stain birthmarks and tattoos). Infrared tomography (IRT) uses a fast infrared focal plane array (IR-FPA) to detect temperature rises in a substrate induced by pulsed radiation. A pulsed laser is used to produce transient heating of an object. The temperature rise, due to the optical absorption of the pulsed laser light, creates an increase in infrared emission which is measured by the IR-FPA. Although the application of IRT to image subsurface cracks due to metal fatigue is a topic of great interest in the aircraft industry, the application to image subsurface chromophores in biological materials is novel. We present an image recovery method based on a constrained conjugate gradient algorithm that has obtained the first ever high quality images of port wine stain blood vessels.

The IRT integral equation relates the measured time sequence of infrared emission frames, $DM_{cp}(x,y,t)$, to the initial space-dependent temperature increase, $DT_{cp}(x,h,z,t=0)$, in individual chromophores. The important physical processes for IRT are represented in a multidimensional integral by the thermal point spread function (PSF), KT(x-x,y-h,z,t), and the camera PSF, $K_C(x-x,y-h)$:

$$\Delta M_{cp}(x,y,t) = \iint_{x',y'} K_c(x-x',y-y')dx'\,dy' \cdot$$

$$\cdot \iiint_{\xi,\eta,\zeta} K_T(x'-\xi,y'-\eta,\zeta,t)\Delta T_{cp}(\xi,\eta,\zeta;t=0)d\xi d\eta d\zeta$$

or, in operator notations:

$$\Delta M_{cp} = K_c * K_T * \Delta T_{cp}$$

The character of the inverse problem is uniquely determined by the thermal PSF K_T, and the camera PSF K_C. The thermal PSF is a separable function of its lateral component K_z and its longitudinal component K_r so that $K_T=K_zK_r$. The lateral kernel is a spatially invariant function of the lateral coordinates (x,y), so removing its effects is a deconvolution problem. However, the longitudinal kernel is a spatially variant function of the longitudinal coordinate z, and it is also a function of time. Consequently, solving the longitudinal problem is computationally more intensive. Furthermore, both problems are highly ill-posed; in particular, the longitudinal kernel has only a few significant singular values.

We have adapted and modified the algorithm described in [1] to solve both the lateral deconvolution problem [2] and the longitudinal inversion problem [3]. These are key components on the way to a more general tomographic reconstruction algorithm and allow us to determine, respectively, the diameter and depth of individual chromophores. The application of this work in a clinical setting at the Beckman Laser Institute and Medical Clinic at UC Irvine is described in Ref. [4], where it is argued that information about diameter and depth is crucial for determining the best laser intensity and pulsewidth for optimum treatment. An existing CRADA between Beckman and LLNL supports the development of a variable pulse width laser that will make this optimum treatment possible.

[1]. "On Applying the Conjugate Gradient Algorithm to Image Processing Problems," Goodman, Johansson, and Lawrence, Chapter 11 in "*Multivariate Analysis: Future Directions*," Elsevier Science Publishers, Amsterdam, 1993.

[2]. "Inversion of Pulsed Photothermal Radiometry Signals in *in vitro* Model Systems," Milner, Goodman, and Nelson. To be submitted to *Applied Optics*, 1994.

[3]. "An Inverse Problem in Pulsed Photothermal Radiometry," Milner, Goodman, Tanenbaum, and Nelson. Submitted to *J. Optical Society of America-A*, 1994.

[4]. "Infrared Tomography for Diagnostic Imaging of Port Wine Stain Blood Vessels," Nelson, Milner, and Goodman. Submitted to *The Journal of Investigative Dermatology, 1994.*

Classification of Heart Valve Condition Using Acoustic Measurements

Gregory Clark

Prosthetic heart valves and the many great strides in valve design have been responsible for extending the life spans of many people with serious heart conditions. Even though the prosthetic valves are extremely reliable, they are eventually susceptible to long-term fatigue and structural failure effects expected from mechanical devices operating over long periods of time. The purpose of our work is to classify the condition of in vivo Bjork-Shiley Convexo-Concave (BSCC) heart valves by processing acoustic measurements of heart valve sounds. The structural failures of interest for Bscc valves is called single leg separation (SLS). SLS can occur if the outlet strut cracks and separates from the main structure of the valve. We measure acoustic opening and closing sounds (waveforms) using high sensitivity contact microphones on the patient's thorax. For our analysis, we focus our processing and classification efforts on the opening sounds because they yield direct information about outlet strut condition with minimal distortion caused by energy radiated from the valve disc.

Our heart valve analysis system consists of algorithms for data acquisition, signal processing, and signal classification. Data acquisition and signal processing are discussed in other papers. This paper concentrates on classification algorithms and results, with brief descriptions of the acquisition and signal processing provided as necessary to understand the inputs to the classifiers. We describe the extraction and selection of spectral features from the spectral estimates computed by the signal processing system. We show that a two-stage classifier, which first classifies individual heart beats and then classifies valves, is very effective in correctly in classifying the valves represented in our data set. We measure performance by constructing confusion matrices, receiver operating characteristic (ROC) curves showing probability of detection and probability of false alarm (sensitivity and 1-specificity), and specifying statistical confidence intervals on the probability of correct classification. We show that even given a very small number of training samples, the classifier achieves probability of correct classification of 100%.

Segmentation and Measurement of Collagen Fibers for Shoulder and Joint Therapy Studies

Laura N. Mascio

Background and Motivation Various shoulder instabilities are debilitating, especially in individuals who perform overhead activities. Physical therapy, open surgical techniques and arthroscopic techniques either do not achieve adequate stability or else severely limit function, so that patients are not able to return to their previous level of activity or performance after treatment. Thermal modification of soft tissues in joints may allow precise alteration of these tissues' mechanical and/or structural properties to enhance joint function without inducing cell death or an inflammatory response. Several studies have evaluated laser energy for tissue welding. An in vitro study of rabbit revealed that at lower energy densities, significant capsular reduction can be achieved without significant alteration of the tissue's mechanical properties. A multi-institutional human study showed patients treated with this technique improved dramatically after surgery, with concomitant increased stability, decreased pain, and enhancement of function, allowing individuals to return to their previous level of activity. The collective findings of laser welding studies are promising, and the next step is to identify the mechanisms responsible for laser-induced capsular tissue alteration, and the short- and long-term effects of non-ablative laser energy on joint capsular tissue. One step toward this goal is to compare the effect of three laser energy densities (selected from the in vitro study) on the histologic properties of the tissue evaluating the architecture of the collagen (including density, fibril diameter distribution, and interfibrillar space) in sheep at various time intervals after surgery. The specific computer algorithms that are being used to make these measurements will be described. Segmentation and Measurement Isolating collagen fibers from their background in digitized electron micrograph images is particularly challenging because of the large variation in size and intensity of the objects. The relative diameters of these fiber cross-sections can range from 4 pixels to 50 pixels in one image, and even more from image to image. Also, the intensity distribution within an object can be highly irregular, especially in fibers that are swollen and misshapen. A segmentation approach, called BackPath, has been developed which exploits the fact that the collagen fibers (in cross-section) are very close together and touching. This may not seem like an advantage, but consider detection of the background instead of the foreground. The BackPath algorithm is designed to trace maximum-cost paths from the top to the bottom of an image, so that each trace must snake through the background, between objects in order to reach its goal. When multiple paths are traced vertically as well as horizontally, the image is divided into irregularly shaped, context-based regions containing one or more whole objects. Each region can be subsequently analyzed to segment the object or objects within them. This concept is applied iteratively within regions and is combined with other methods to further segment the individual objects in the image. Intermediate results yield images which are suitable for measuring various features. For example, the average intensity under a trace through the background gives a good estimate for a threshold value that will segment the background. This threshold fills in patches of background and extra space around objects that would not be filled by the traces alone. The sum of the threshold result and all of the vertical and horizontal traces yields a good estimate for the amount of space in the image that is occupied by background vs. objects, although touching objects may remain connected in this result. While this measurement is appealing to the human eye, changes in the threshold value can alter the resulting area ratios and yet be equally visually appealing. The effect and extent of this variability will be examined. Final identification of each fiber and its extent will also be discussed

Dental Diagnostics Using Optical Coherence Techniques

H. Nathel[†‡], B. Colston[†∨], G. Armitage[‡], and L. Otis[•]

[†]Lawrence Livermore National Laboratory
[‡]Department of Stomatology, U. C. San Francisco
[∨]Department of Biomedical Engineering, U. C. Davis
[•]Department of Oral Diagnosis, University of Connecticut

Optical radiation can be used for diagnostic purposes in oral medicine. However, due to the turbid, amorphous, and inhomogeneous nature of dental tissue conventional techniques used to transilluminate materials are not well suited to dental tissues. Optical coherence techniques either in the time- or frequency-domain offer the capabilities of discriminating scattered from unscattered light, thus allowing for imaging through turbid tissue. Currently, using optical time-domain reflectometry we are able to discriminate specular from diffuse reflections occurring at tissue boundaries. We have determined the specular reflectivity of enamel and dentin to be approximately 6.6 x 10^{-5} and 1.3 x 10^{-6}, respectively. Implications to periodontal imaging will be discussed.

This work was performed under the auspices of the U. S. Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-ENG-48 and support for this research was provided from Grant 1 RO1 DE11154-01 from the National Institute of Dental Research.

VISION/IMAGE DISPLAY

Diffractive Optics for Human Vision Correction

Donald Sweeney

Abstract not available.

Real-Time Vision Systems

Robert K. Johnson, Jose E. Hernandez, Shin-yee Lu

Introduction

Many industrial and defense applications require an ability to make instantaneous decisions based on sensor input of a time varying process. Such systems are referred to as "real-time systems" because they process and act on data as it occurs in time. When a vision sensor is used in a real-time system, the processing demands can be quite substantial, with typical data rates of 10-20 million samples per second.

A real-time Machine Vision Laboratory (MVL) was established in FY94 to extend our years of experience in developing computer vision algorithms to include the development and implementation of real-time vision systems. The laboratory is equipped with a variety of hardware components, including Datacube image acquisition and processing boards, a Sun workstation, and several different types of CCD cameras, including monochrome and color area cameras and analog and digital line-scan cameras. The equipment is reconfigurable for prototyping different applications. This facility has been used to support several programs at LLNL, including O Division's Peacemaker and Deadeye Projects as well as the CRADA with the U.S. Textile Industry, CAFE (Computer Aided Fabric Inspection). To date, we have successfully demonstrated several real-time applications: bullet tracking, stereo tracking and ranging, and web inspection. This work has been documented in the ongoing development of a real-time software library.

Real-time Machine Vision Library

We have developed an image processing library, RTVLIB, to facilitate and accelerate the software development of real-time image processing applications. The library was written on top of the Imageflow software used to program Datacube hardware. The library is composed of a low level and a high level layer. The low level layer contains functions to initialize the connections of the different hardware processing elements, which is usually tedious and extremely prone to error. An example is a function to initialize a particular device to add two images. The high level layer contains functions for creating the event driven processing modules, known as PATs (Pipeline Altering Threads), which are responsible for performing the desired image processing operations in real-time. An example is a PAT that waits for the next frame to be available and displays a frame difference in real-time. RTVLIB was developed and used with the applications described below.

Bullet Tracking

Our collaboration with the Deadeye Project in O-Division resulted in the world's first live demonstration of detecting and tracking bullets as they fly through the air. We developed a real-time image processing system that is capable of processing over 200 images per second from a specialized sensor to detect and track moving bullets. Bullet trajectories are formed and displayed in real-time as they occur. In fact, during testing at a firing range, we were able to demonstrate that the trajectories were displayed before the bullet actually reached the target.

The system that we delivered, known as Lifeguard, is capable of operating in one of four modes as desired: (1) continuously monitor the scene and detect and display trajectories, (2) archive all images to a real-time disk without processing, (3) simultaneously process and store data, and (4) process data from previously stored files. This technology will be used to develop monitoring devices capable of locating and responding to sources of gunfire. The response would vary depending on the environment, civilian or military, but could range from recording a picture of the gunman to firing back. Lifeguard has received international news coverage, and it is expected to develop into a major program in FY95. A video of several news clips demonstrating the system will be shown.

Stereo Tracking

A real-time stereo tracking system was developed and implemented. The system is capable of tracking an object feature point from a pair of synchronized CCD cameras and providing an estimate of the corresponding three-dimensional (3D) coordinates at the full frame rate of the cameras (30 Hz). An operator first selects an object of interest in the field of view of one of the cameras using a graphical user interface. The system then begins tracking the selected object and displays the 3D coordinates of the object's centroid as the default feature point. The operator can change the feature point at will by simply clicking on any point in the object. The feature point is tracked independent to changes in the object's scale, translation, and rotation.

The system has been demonstrated as a proof of concept at MVL using a robot to move an object through various trajectories. The current application for this capability is for developing extremely precise targeting devices for the Peacemaker Project. One proposed device would consist

of an unmanned remote control air vehicle equipped with a stereo tracking system. This vehicle could be flown into hostile environments without endangering the lives of a pilot and crew. The stereo tracking system could be used to precisely locate and disable strategic targets which would also minimize enemy casualties.

Web Inspection

We are currently developing a web inspection system in support of the CRADA with the American Textile Industry, CAFE. The current system consists of a machine for moving fabric or similar materials, such as paper, at speeds up to 80 yards per minute, a line-scan camera, and a high frequency aperture fluorescent lamp. Real-time computer vision software has been developed and implemented to detect printing defects. The goal of this project is to create a real-time inspection system which can be placed directly on-line in a manufacturing plant in order to isolate defects and improve efficiency. The prototype system was recently demonstrated at the Bobbin Show, which is one of the largest textile trade shows in the U.S. It was also shown on a nationally televised news program. A video of the news clip will be shown. We expect this work to generate new opportunities in other web inspection applications as well as expand our expertise in the development of real-time vision systems.

Nonlaser-Based 3D Surface Imaging

Shin-yee Lu, Robert K. Johnson, Robert J. Sherwood

3D surface imaging refers to methods that generate a 3D surface representation of objects or a scene under viewing. Laser-based 3D surface imaging systems are commonly used in manufacturing, robotics and biomedical research.

Although laser-based systems provide satisfactory solutions for most applications, there are situations where non laser-based approaches are preferred. The issues that make alternative methods sometimes more attractive are: (1) real-time data capturing, (2) eye-safety, (3) portability, and (4) work distance. The focus of this presentation is on generating a 3D surface from multiple 2D projected images using CCD cameras, without a laser light source. Two methods are presented: stereo vision and depth-from-focus. Their applications are described.

Stereo Vision

Stereo vision uses two cameras mounted at a fixed base line and focused

on the same object. One image is captured from each of the two cameras simultaneously. This stereo pair of images is registered pixel to pixel. Different perspectives of the cameras and depth variation on the object can cause occlusion as well as scaling changes on the two images. The accuracy of 3D reconstruction depends on the stereo registration algorithms that match image intensity and intensity changes of two images. An active light source can be used to project light patterns on the object and create sharp intensity changes for registration.

Stereo vision is an ideal method for 3D imaging of the human body and movement for two reasons: real-time data capturing and eye-safety. Examples of 3D human face and body reconstruction will be presented.

Depth-from-Focus 3D Microscopic Imaging

The shallow depth of focus of an optic system can also be explored to generate 3D information of an object. The depth-from-focus approach uses image processing techniques to deduce object depth information from optical sections. Optical sections are a series of images obtained from focusing on an object at different depths. The advantage is that it does not require stereo imaging heads and image registration software, nor does it require scanning motion like a confocal microscope system. We will describe a system of low cost microscopic imaging hardware that is used to obtain optical sections, and the depth-from-focus computer algorithm that generates the 3D reconstructed object.

The system can provide better than 1 micron resolution on a focal plane and 10 micron resolution along the optic axis. It is being applied to micro-assembly automation.

Scientific Graphical Displays on the Macintosh

Stanley Grotch

Abstract

In many organizations scientists have ready access to more than one computer, often both a workstation (e.g., SUN, HP, SGI) as well as a Macintosh or other PC. The scientist commonly uses the work station for "number-crunching" and data analysis whereas the Macintosh is relegated to either word processing or serves as a "dumb terminal" to a larger main-frame computer. In an informal poll of my colleagues, very few of them used their Macintoshes for either statistical analysis or for graphical data display.

I believe that this state of affairs is particularly unfortunate because over the last few years both the computational capability, and even more so, the software availability for the Macintosh have become quite formidable. In some instances, very powerful tools are now available on the Macintosh that may not exist (or be far too costly) on the so-called "high end" workstations. Many scientists are simply unaware of the wealth of extremely useful, "off-the-shelf" software that already exists on the Macintosh for scientific graphical and statistical analysis.

This paper is a very personal view illustrating several such software packages that have proved valuable in the author's own work in the analysis and display of climatic datasets. It is not meant to be either an all-inclusive enumeration, nor is it to be taken as an endorsement of these products as the "best" of their class. Rather, it has been found, through extensive use that these few packages were generally capable of satisfying my particular needs for both statistical analysis and graphical data display. In the limited space available, the focus will be on some of the more novel features found to be of value.

The following discussion is divided into three sections, the first two illustrating Macintosh software for statistical data analysis and for graphical data display. The final section will summarize the work and offer some comments regarding the future.

IMAGING HARDWARE

Synthetic Aperture Radar Capabilities in Development

Mark Miller

The Imaging and Detection Program (IDP) within Laser Program is currently developing an X-band Synthetic Aperture Radar (SAR) to support the Joint US/UK Radar Ocean Imaging Program. The radar system will be mounted in the program's Airborne Experimental Test-Bed (AETB), where the initial mission is to image ocean surfaces and better understand the physics of low grazing angle backscatter. The large bandwidth system in conjunction with high quality motion data will provide high resolution in both range and azimuth. The AETB, a medium sized jet aircraft, will have the versatility to host a variety of SAR applications due its good motion stability and the capability of attaining high altitudes and medium ranges. The Synthetic Aperture Radar presentation will discuss its overall functionality and a brief discussion on the AETB's capabilities. Vital subsystems including radar, computer, navigation, antenna stabilization, and SAR focusing algorithms will be examined in more detail.

The first application of the AETB/SAR will be to image ocean-like surfaces to study naturally occurring and man-made internal waves. SAR's have been successfully used in several radar ocean imaging projects, where they were primarily used to image water surfaces at 20 to 30 degrees incident. IDP has fielded several real aperture radar (RAR) systems which were located on hillsides and used to image water bodies at a five degree grazing angle. At low gazing angles (LGA), these radars have been able to image much better than predicted. The new scattering mechanisms apparently responsible for this increased performance require addition experiments. IDP will use this SAR capability to study this scattering with hopes to create a better imaging system.

Many sensor packages have been used to study internal waves. For this reason, the AETB will be "generically" reconfigured to host a variety of RF and optical hardware. The test-bed configuration will include: a long narrow belly radome, several equipment racks and experimental hardware mounting plates, four experiment operation/ control stations, and conditioned electrical power in the form of 60 and 400Hz. In the SAR configuration, the antenna will be mounted in the radome and the remaining hardware will be mounted in the equipment racks and consoles located throughout the pressurized fuselage.

To better accommodate multiple missions that require different sensor packages, all of the experimental hardware that make up the SAR system has been designed to be modularized and easily reconfigured. The major building blocks consists of a coherent S through Ku band (2-18 GHz) radar system which will initially be configured to operate in X band (9-10 GHz), a parallel computer system that provides real-time SAR processing and analysis, a navigation system to accurately determine antenna motion and flight path information, and a stabilized antenna gimbal to eliminate the aircraft's axial motion as seen by the antenna and precisely point the antenna to specific coordinates.

The key component of the SAR is the radar unit itself, which is a very versatile system because of its broad frequency band of operation, numerous waveforms, high power, high bandwidth, variable pulsewidths and variable PRF capabilities. The S, C, X, and Ku band system uses direct digital synthesis capable of generating an infinite number of waveforms. The most practical waveforms for our applications are the linear frequency modulated (LFM) waveform with up to 125 MHz instantaneous bandwidth and the stepped LFM waveform with up to 500 MHz bandwidth. The coherent radar has an Inphase and Quadrature demodulator which will be sampled at rates up to 200 MHz with an 8 bit A/D. Large range swaths can easily be achieved even in the highest resolution mode due the system's 4096 range cells. An 8KW X-band Traveling Wave Tube Amplifier (TWTA) will provide sufficient power to image sea surface at low grazing angles as well as other targets with small cross-sectional areas or targets located at long distances.

A Mercury i860/Sparc based computer system controls all facets of the radar and peripheral systems and more importantly has the power to perform real-time SAR processing. The ten i860 processors that make up the data processor provide enough computing power to process 1024 range cells at a pulse repetition frequency of 1 KHz, including motion compensation. With the future implementation of dedicated FFT hardware the data processor will be able to process even larger SAR images with increasing resolution. Additional image processing for SAR analysis will also be provided in the initial system.

To accommodate the high data rates of the radar system and the data processor, a high speed digital tape recorder will be used to archive the radar data and system parameters. The recorder can store up to 48 GBytes of data on a cartridge at rates up to 30 MBytes/sec. A disk array system will provide supplemental recording capabilities, i.e., record processed SAR images, and can also serve as a back up to the tape recorder.

Another subsystem is the stabilized platform which will control the antenna pointing vector and allow mapping from either side of the aircraft from the horizontal to the nadir position. Roll and yaw will be stabilized to within a few tenths of one degree; no active pitch control will initially be implemented although the gimbal will be mounted such that the mean pitch angle of the antenna will be zero degrees.

The high resolution antenna motion and position information which are required by three of the subsystems will be provided by a Global Positioning System aided Inertial Navigation System (GPS/INS). GPS is a satellite based system that provides stable long term accuracies but with only one second updates, while the INS provides good short term stability but suffers long term drift errors. When implemented as system, they prove to be an excellent motion sensor by using the GPS data to model and correct the INS errors. The SAR processor requires the navigation data to perform motion compensation and properly focus SAR images. A course correction display will also use this data to aid the pilot in maintaining the proper flight path. In addition, all of the axial motion and position information will be used by the gimbal to stabilize the antenna.

The high resolution real-time imaging SAR is being developed for ocean imaging but offers many features required by other applications. Built in modularity will simplify reconfigurations and minimize down-time between experiments. Addition versatility is also provided by the high performance AETB which has been "generically" modified to accommodate various hardware configurations. When the AETB/SAR completes its engineering test in March and has proven its high resolution capability, additional operations will be developed to fully utilize this SAR asset.

Advanced CCD Camera Developments

Alan Condor

We will introduce and discuss two charge coupled device (CCD) camera systems that we have recently developed, describing briefly the hardware involved, and the data obtained in their various applications. These systems include the 400 frame-per-second 512 x 512 CCD camera, and a compact, 16-bit scientific grade 1024 x 1024 CCD camera.

The Advanced Development Group Defense Sciences Engineering Division has been actively designing, manufacturing, fielding state-of-the-art CCD camera systems for over a decade. We originally developed these systems for the nuclear test program to record data from underground nuclear tests. Today, new and interesting applications for these systems have surfaced. We are continuing in the development of advanced CCD camera systems, with the new CCD camera that will allow a experimenters to replace film for x-ray imaging at the JANUS, USP, and NOVA laser facilities.

We developed a high speed 512 x 512 CCD camera, capable of framing 400 digital images per second with a 9-bit dynamic range for the nuclear test program. This CCD camera has since found use in imaging projectiles launched from the SHARP gun, and when outfitted with an IR (1.0 μm to 1.6 μm) sensitive proximity focused diode, has obtained images of shaped charge explosions. This CCD camera is unique, not only because of its fast frame rate, but also for its 1.44 Gb/s fiber optic data link and recording system, specifically developed for this camera, making remote, high speed, high resolution, real time imaging possible.

We developed a very compact (60 x 60 x75 mm^3), vacuum compatible, large format (25 x 25 mm^2, 1024 x 1024 pixels) CCD camera for digital imaging of visible and ultraviolet radiation, soft to penetrating x-rays (≤ 20 keV), and charged particles. This camera provides a suitable replacement for film with a linear response, dynamic range and intrinsic signal-to-noise response superior than current x-ray film, and provides real-time access to the data. The spatial resolution for the camera (<25 μm) is similar to typical digitization slit or step sizes used in processing film data. This new large format CCD camera has immediate applications as the recording device for streak cameras or gated microchannel plate diagnostics, or when used directly as the detector for x-ray, xuv, or optical signals. This is especially important in studying high-energy plasmas produced in pulse-power, ICF, and high-powered laser-plasma experiments, as well as other medical and industrial applications.

Clementine Sensor Suite

Arno G. Ledebuhr, Robert E. Priest,
Joseph F. Kordas, Noel R. Sewall,
Isabella T. Lewis, Hye-Sook Park,
Robert F. Hills, Diana E. Sackett,
Michael J. Shannon and Lyn D. Pleasance

Introduction

Lawrence Livermore National Laboratory (LLNL) designed and built the suite of six miniaturized light-weight space-qualified sensors utilized in the Clementine mission. A major goal of the Clementine program was to demonstrate technologies originally developed for Ballistic Missile Defense Organization (BMDO) programs. Modifications were made to the sensors to fine tune their spectral ranges and responsibilities to gather the highest quality and greatest quantity of science data at the moon utilizing the basic BMDO sensor design forms. These mission instrument suite consisted of seven sensors. A pair of WFOV Star Tracker cameras used to provide attitude determination as part of the spacecraft ACS and were also used for WFOV lunar and earth imaging. The main imaging sensor complement included a UV/Visible (UV/VIS) Camera, a High-Resolution Imaging (HiRes) Camera, a Near-Infrared (NIR) Camera, and a Longwave Infrared (LWIR) Camera. The HiRes was the imaging portion of an LIDAR system which also incorporated an advanced solid state laser transmitter and a laser ranging channel. The sensors viewed reflected solar and earth illumination off the lunar surface, while the LIDAR system gathered lunar topography data by measuring the altitude of the spacecraft via the time delay for the return of a pulsed laser beam. Automated pre-flight calibration measurements at LLNL provided a comprehensive data set for calibration of the lunar science data. State-of-the-art data compression techniques maximized the quantity of data gathered at the lunar surface. The Clementine mission collected over 1.7 million images that included a multi-spectral map of the entire lunar surface in 11 spectral bands and a topographical map of the moon to 40 m vertical resolution. This data is currently be analyzed and is being archived in NASA's Planetary Data System.

We will present an overview of each of these sensors and some preliminary on-orbit performance estimates. The basic subsystems of these sensors include optical baffles to reject off-axis stray light, light-weight ruggedized optical systems, filter wheel assemblies, radiation tolerant focal plane arrays, radiation hardened control and readout electronics and low mass and power mechanical cryogenic coolers for the infrared sensors. All sensor systems were space qualified to MIL-STD 1540B. The sensors have a flexible set of control parameters such as integration time, gain and offset, that allowed for the collection of a high quality set of science data and to meet a wide range of science objectives. Descriptions of each sensor type will be given, along with design specifications, photographs and on-orbit data collected.

Wide-Field-of-View (WFOV) Star Tracker

The WFOV Star Tracker consists of a miniaturized, 290 gram, CCD imaging camera. The Star Tracker camera is unique among spacecraft Star Tracker systems in that the camera has a wide-field-of-view (29° x 43°) provided by a fast (F/1.25) concentric lens design. A fiber optic field-flattener is used to couple the lenses curved image surface to a flat CCD array. The CCD used is manufactured by Thomson-CSF and has a 576 x 384 pixel format. The pixel size of this array is 23 x 23 microns and the spectral response is from 0.4 - 1.1 microns. The primary function of the Star Tracker camera is to provide stellar images that are processed against an on-board star catalog that establishes the absolute angular reference for navigation. When operated with its robust star matching algorithm running on the sensor processor, this system thus forms a Stellar Compass system. The Star Tracker provides inertial reference attitude for a spacecraft to demonstrated accuracies of 150 microradians in pitch and yaw and 450 microradian in roll. The extremely robust star matching algorithm can generate attitude quaternions in less than 100 ms on a RISC-3000 processor. The sensor system matches stars as dim as magnitude Mv 4.5 in 4-pi steradians of the stellar sky. This system solves the lost in space problem on a single image without any a prior knowledge of orientation with respect to the celestial sphere.

UV/Visible (UV/VIS) Camera

The UV/Visible camera is a miniaturized multi-spectral imaging sensor massing only 430 grams. This space-proven camera provides reliable solid-state imaging in the ultra-violet (300 nm) to near-infrared (1100 nm) regime. The sensor uses a catadioptric lens with fused silica lenses. The lens aperture has a 46 mm diameter with a 90 mm focal length and with a corresponding speed of F/1.96. The sensor's FOV is 4.2 x 5.6 degrees and the pixel IFOV is 255 microradians. The CCD array is Thomson's TH7863 frame-transfer device which uses frame-transfer electronic shuttering and is a sister to the Star Tracker's TH7883. This shuttering is accomplished by rapidly shifting the active pixel area into an identically sized on-chip and light-shielded storage area. This clears the active region and the clocking is then paused for the (13-bit programmable) integration time. Following the integration period the captured image is rapidly shifted into this storage area from which the image is read out. Thus the array has only half of the imaging pixels (288 x 384) as does the Star Tracker. The pixel size is 23 x 23 microns square. The CCD is overcoated with a metachrome phosphor which converts UV radiation into visible wavelengths and extends the quantum efficiency of the CCD below 400 nm. Six (6) spectral bands are selectable from a filter wheel which is controlled through the serial-addressable synchronous interface (SASI). Six filter positions are available using a space-rated 90° No. 5 stepper motor geared at 6:1 to the filter wheel, a series of Hall Effect sensors for position determination, and a small electronics card for command and operation. This system is modular and functionally separate from the imaging camera. During the Clementine mission the six spectral bands provided radiometric mineral typing data for the Earth's moon.

LIDAR System

LLNL developed a space-qualified LIDAR (Light Detection And Ranging) system for use on Clementine. The LIDAR receiver system consists of a High Resolution (HiRes) Imaging channel which incorporates a custom intensified multi-spectral visible camera combined using a beam splitter assembly with a Laser Ranging channel which uses an avalanche photo-diode (APD) for laser pulse timing. These receiver channels use a common Cassegrain telescope assembly and are bore-sighted with a light-weight McDonnell-Douglas diode-pumped Nd:YAG laser transmitter. The telescope has a 13.1 cm aperture and a 1.25 m focal length and a resulting speed of F/9.5. The imaging FOV is 0.3 x 0.4 degrees and has the same pixel count and size as the UV/Vis camera. However with the long focal length telescope the pixel IFOV is 18 microradians. This sensor uses the SASI electrical control interface and matches that of the Star Tracker and UV/Vis camera. During the lunar mapping portion of Clementine, the Laser ranging system produced remarkable altimetry data from as far as 640 km altitudes, craters up to 12 km deep were discovered during the Clementine LIDAR mapping, far deeper than previously known. The imaging channel produced high resolution multi-spectral images for radiometric mineral typing of the lunar surface.

Clementine Near-Infrared (NIR) Camera

This miniaturized 1.9 kg, space-proven near-Infrared (NIR) Camera provides reliable solid-state imaging from 1.0 to 3.0 μm. The camera uses a state-of-the-art Amber 256 x 256 Indium Antimonide (InSb) FPA cryogenically cooled to 70K operating temperature by a Ricor K506B Stirling cycle cryocooler. The FPA operated at 70 ± 0.5 K in operation, and showed excellent stability over the more than 500 hours of operation in space. The optical system was an catadioptic lens with a 29 mm aperture a 96 mm focal length and a speed of F/3.33. The pixel IFOV is 400 microradians and the pixel size is 38 x 38 microns. The optics design incorporates an internal relay of ZeSn refractive elements with 100% cold shield efficiency. Wavelength range was controlled with an external warm filter wheel with 6 bands between 1.0 to 3.0 microns. The camera's 8 bit readout electronics had programmable features including four integration times, 5 bits of gain, and 8 bits of offset. Offset is subtracted before gain is applied with 0 V to full well range selectable in 1/255 full well increments. On-board DC-DC converters provide conditioned power forms to the camera allowing it to operate off raw 30 ± 6 V.

Clementine Longwave Infrared (LWIR) Camera

This miniaturized 2.1 kg, space-proven Longwave Infrared (LWIR) Camera provides reliable solid-state imaging from 8.0 to 9.5 μm. The camera uses a state-of-the-art Amber 128 x 128 Mercury Cadmium Telluride (HgCdTe) FPA cryogenically cooled to 65K operating temperature by a Ricor K506B Stirling cycle cryocooler. The FPA operated at 65 ± 0.5 K in operation, and also showed excellent stability over the more than 500 hours of operation in space. The telescope has a 13.1 cm aperture a 35 cm focal length and a speed of F/2.67. The pixel IFOV is 140 microradians and the pixel size is 50 x 50 microns. The optics design incorporates an internal relay of ZeSn

refractive elements with 100% cold shield efficiency. Wavelength range was controlled by the cold filter to 8.0 to 9.5 microns. The Clementine LWIR Camera electronic design is virtually identical to the Clementine NIR camera, with a minor alteration in line readout to compensate for the difference in column count between the two arrays (256 for the NIR; 128 for the LWIR). The NIR and LWIR cameras also share a common cryocooler and dewar design, with minor modifications made to accommodate cold shield and cold filter differences.

Clementine Auto Exposure Control Software

Ernest Arnold

Overview

The primary mission of the Clementine program was to test technology developed under the auspices of BMDO (the Ballistic Missile Defense Organization). A secondary goal of the program was to provide astronomical data to the scientific and educational community. The mission plan developed to accomplish these goals included complete mapping of the lunar surface and a close fly-by of a near-Earth asteroid, 1620 Geographos.

Exposure control for the Clementine mission was driven by mission phase requirements and sensor characteristics. Thus, there were a total of twelve algorithms developed for three primary mission phases and the four imaging sensors (two additional sensors operated as star trackers). The three mission phases in question were lunar mapping, distant observation of the asteroid for the purpose of tracking, and close-up viewing (as close as 100 Km) of Geographos. The four non-star tracker sensors consisted of an UltraViolet/Visible (UV/Vis) camera, a High Resolution (HiRes) camera with a built-in LIDAR (Light Detection And Ranging) unit, a Near InfraRed (NIR) camera, and a Long Wave InfraRed (LWIR) camera. Due to lack of test time and uncertainties about the imaging environment, numerous input parameters were provided in the algorithms to allow extensive tuning of the exposure control during the mission.

Camera Characteristics Relating to Exposure Control:

The exposure controls for each camera were somewhat different, requiring an unique algorithm for each camera and each class of algorithm.

Lunar Exposure Control:

The lunar exposure control relies upon on-board knowledge of spacecraft position, attitude, ground speed, distance from the lunar surface, date, and time as well as characteristics of the sensors to accomplish exposure control. A more sophisticated algorithm would include reference to a database of terrain effects. Lens and focal plane array temperatures were considered by algorithms used for the two infrared cameras.

Three important angles are supplied to the algorithm from the image processing task, IPEXEC. The solar phase angle is the angle having as its vertex the center of the moon and extending in one direction to the spacecraft and to the sun in the other direction. The solar incident angle runs from the normal to the point of observation to the center of the sun. The emission angle describes the angle between the normal to the point of observation and the vector extended toward the spacecraft. Range (distance from the spacecraft to the observation point), ground speed, and the three angles were calculated by IPEXEC, given the distance between the spacecraft and the moon's center and the unit vector in the line of sight direction.

A function approximating cosine (latitude) is calculated from the three angles. The function also responds to the effect of off-nadir observations. The goal for exposure control is to fill the A/D well with the scene maximum - scene minimum + some margin (to prevent clipping at either extreme). The maximum integration time is chosen to minimize pixel-smearing at current altitude and relative ground speed. $Signal_{scene\_max}$ is calculated as $scaling\_factor_{uploadable} \bullet integration\_time_{max}$.

Point Source Exposure Control:

The objective of the point source class of algorithm is to control camera exposure such that a target can be readily detected, using a centroided N-brightest pixel detection scheme. This is accomplished by attempting to control the intensity of the ten brightest pixels of the next image to an optimal gray level, based on the results of previous images.

The notion of camera exposure gain is maintained. For example, the exposure gain for a given filter wheel position of the UV/Vis camera is:

exposure_gain = filter_transmissivity • integration_time • video_gain

This value is maintained as a running average of the last five exposure adjustments for a given filter. The integration time and video gain settings are then derived from this exposure gain value. The new value of integration time, for example, is calculated as:

integration_time = exposure_gain / (next_filter_transmissivity • optimal_video_gain)

A different filter may be used for the next image, so its transmissivity must be used, not that of the image just acquired (unless they are the same). A preferred video gain setting is selected unless the result puts the integration time out of a desirable range. If that is the case, the video gain is set from the exposure gain and the acceptable value of integration time closest to the value that was calculated.

Offset adjustment is only applied to compensate for effects due to changes in the video gain setting.

Extended Scene Exposure Control:

The extended scene class of exposure control software was intended for the brief period of time when the spacecraft was to be close to the asteroid. Image feedback was necessary since we know little about the asteroid's size and geometry. The goal during the near-asteroid imaging was to quickly converge the algorithm while providing maximum gray level information. The lunar exposure control software was utilized to adjust exposure control parameters. A front end was provided that essentially replaced the sun-moon-spacecraft angle mechanism to that of image feedback through evaluation of the resultant histogram. Some percentage (typically 0.5%) of the total number of pixels in the image are truncated from either end of the histogram. The corresponding gray levels are returned as ctsMin and ctsMax. The intent is to have ctsMin to be close but not equal to zero and ctsMax to be close but not equal to 255. This ensures that the scene is being matched well to the camera A/D range and is therefore being well characterized. The next step is to calculate sceneMin and sceneMax value from ctsMin and ctsMax. For the UV/Vis and HiRes cameras:

sceneMin = (((ctsMin + offset / offset_factor - 51) • electrons_per_bit) / gain)
/ (integration_time • scaling_factor)

sceneMax = (((ctsMax + offset / offset_factor - 51) • electrons_per_bit) / gain)
/ (integration_time • scaling_factor)

where integration_time is the integration time in seconds
offset is the camera offset level
offset_factor is the number of gray levels corresponding to one offset level (nominally 8)

electrons_per_bit is one of 1000, 350, or 150e⁻/bit (video gain setting such that this many photon-induced electrons are required to change the A/D setting by one bit) gain is the exposure gain due to the image intensifier gain setting

Other values are calculated on a per sensor basis in order to accommodate required inputs for the algorithm designed initially for lunar exposure control.

The values of ctsMin and ctsMax are adjusted as necessary to prevent underflow and overflow conditions. If ctsMin is less than 1 or greater than 250, its is adjusted exponentially with respect to the number of attempts to bring the exposure within the proper range.

Test Environments

Testing the exposure control software was very difficult due to a lack of available sensors, the extremely compressed time scale, and the absence of a fully integrated test environment until six weeks prior to launch. Also, the test environment was located in Alexandria, Virginia and my home is in Livermore. Initial testing was accomplished using stand alone codes. The user entered parameters and observed the calculated exposure parameters for the lunar exposure control. The point source and extended scene exposure control algorithms were tested against user-supplied exposure parameters and image files. Eventually, the last two algorithms were integrated into a real-time imaging capture, display, and exposure control program for more realistic testing.

The first integrated testing took place 10/28/93 with three months remaining before launch. This first attempt was made on the spacecraft at the Naval Research Laboratory in Washington, DC and did not even invoke my software. Problems were discovered and addressed during late night spacecraft tests on 10/31/93, 11/11/93, and 12/2/93. The opportunities to test were indeed few and far between. In December, response to my plea brought about a separate, integrated test environment in the control center in Alexandria, across the Potomac.

Each of the algorithms was tested on the in-flight spacecraft, though parameter-tuning was never satisfactorily completed.

Results From In-Flight Tests:

Each of the three classes of algorithm was tested with a measure of success during the three month mission. A decision was made to manually load exposure tables and refine settings during the two-month lunar orbiting phase due to a lack of sufficient testing time. Since the mission was prematurely aborted due to an untimely crash of the main processor and subsequent release of attitude control

fuel, the other two algorithms were not used during the intended phases.

On March 7, 1994, during orbit 73, the most conclusive of the lunar exposure control tests was conducted between 50° and 90° N latitude. The exposure settings were consistent with the anticipated cosine(latitude) function, with two caveats. These anomalies were most likely the result of problems related to software configuration management and operator script over-run.

Two days before the Clementine spacecraft went into a spin, the point source control algorithm was tested for the UV/Vis and NIR cameras. For the UV/Vis camera, the point source control algorithm overshot its target of 200 counts, then undershot and gradually homed in on 200 counts approximately. Parameters could be adjusted in order to converge more quickly, though the results are more than adequate for target detection purposes. Filter 6 was used to drive the exposure control for all filters. The other filters converged at different gray levels due to apparent errors in the filter transmissivity values. These values were easily modifiable by the operator on the ground, using a script.

On the following day, May 6, 1994, the extended scene algorithm was tested. Again, filter 6 was used to control exposure settings for images at all filter settings. After three adjustments to the exposure, the algorithm converged with ctsMin approximately equal to 30 counts and ctsMin at roughly 220 counts.

Mercuric Iodide Room-Temperature Array Detectors for Gamma-Ray Imaging

Bradley E. Patt, Ph.D

Significant progress has been made recently in the development of mercuric iodide detector arrays for gamma-ray imaging, making real the possibility of constructing high-performance small, light-weight, portable gamma-ray imaging systems. New techniques have been applied in detector fabrication and then low noise electronics which have produced pixel arrays with high-energy resolution, high spatial resolution, high gamma stopping efficiency. Measurements of the energy resolution capability have been made on a 19-element protypical array. Pixel energy resolutions of 2.98% fwhm and 3.88% fwhm were obtained at 59 keV (241-Am) and 140-keV (99m-Tc) respectively. The pixel spectra for a 14-element section of the data is shown in the figure (top), together with the composition of the overlapped individual pixel spectra (bottom). These techniques are now being applied to fabricate much larger arrays with thousands of pixels. Extension of these principles to imaging scenarios involving gamma-ray energies up to several hundred keV is also possible. This would enable imaging of the 208 keV and 375-414 keV 239-Pu and 240-Pu structures, as well as the 186 keV line of 235-U.

High Performance Image Processing of SPRINT

Tony DeGroot

This talk will describe computed tomography (CT) reconstruction using filtered back-projection on SPRINT parallel computers. CT is a computationally intensive task, typically requiring several minutes to reconstruct a 512x512 image. SPRINT and other parallel computers can be applied to CT reconstruction to reduce computation time from minutes to seconds. SPRINT is a family of massively parallel computers developed at LLNL. SPRINT-2.5 is a 128-node multiprocessor whose performance can exceed twice that of a Cray-Y/MP. SPRINT-3 will be 10 times faster. Described will be the parallel algorithms for filtered back-projection and their execution on SPRINT parallel computers.

Advanced Wavefront Measurement and Analysis of Laser System Modeling

C. R. Wolfe, J. M. Auerbach

High spatial resolution measurements of the reflected or transmitted wavefronts of large aperture optical components used in high peak power laser systems is now possible. These measurements are produced by phase shifting interferometry. The wavefront data is in the form of 3-D phase maps that reconstruct the wavefront shape. The emphasis of this work is on the characterization of wavefront features in the mid-spatial wavelength range (from 0.1 to 10.0 mm) and has been accomplished for the first time. Wavefront structure from optical components with spatial wavelengths in this range are of concern because of their effects in high peak power laser systems. At high peak power, this phase modulation can convert to large magnitude intensity modulation by non-linear processes. This can lead to optical damage. We have developed software to input the measured phase map data into beam propagation codes in order to model this conversion process. We are analyzing this data to:

1. Characterize the wavefront structure produced by current optical components

2. Refine our understanding of laser system performance

3. Develop a database from which future optical component specifications can be derived

Summary

Modern optical phase shifting interferometry is highly developed. Commercial instruments are currently available with apertures as large as 450 mm (18"). Only static fringe analysis is possible at larger apertures. The interferometer aperture used for measurement and the detector array size determine the spatial bandwidth detectable in the wavefront. Using the phase shifting technique, the phase of each pixel of the solid state detector is modulated and the contour of the wavefront is synthesized. Data can be filtered, if necessary, to emphasize either high or low frequency components. Commercial software allows the analysis of single "line-outs" of wavefront data. This is usually sufficient for evaluating random wavefront structure (i.e. noise). Random noise sources in optical components are produced by uncorrelated grinding or polishing processes, defects, or random index variation of optical materials. On the other hand, sources of regular wavefront structure are also present in optical components and these are of great interest for advanced laser system design and optimization purposes. Regular wavefront structure can arise from index variation of materials (i.e. material striae). It can also originate on surfaces produced by correlated small tool grinding or polishing, or from diamond turning processes. With current software, the detection and analysis of regular wavefront features requires accurate alignment of multiple line-outs with the features in the wavefront. We are beginning to overcome this limitation by extending wavefront analysis to 2-D using advanced image analysis software.

The design and optimization of advanced laser systems requires detailed beam propagation modeling. The new wavefront phase map data containing information about mid-spatial wavelength structure is critical for the design of new high peak power solid state laser systems, such as the NIF and Beamlet. Wavefront structure from optical components is imprinted on the laser beam as phase disturbances. When the mean beam intensities approach values of several GW/cm^2, non-linear processes can convert phase structure into intensity structure on the beam. This intensity structure can grow to large values if it is not eliminated by the laser system's spatial filters. When the magnitude of wavefront structure caused by optical components is too large, the induced modulation of the beam intensity (or fluence) can exceed optical component damage thresholds. In order to prevent this situation, detailed modeling of the laser system performance is done. This modeling can help determine the operating conditions and optical component specifications necessary to maintain good beam quality.

Our goal is to ultimately derive engineering specifications for optical components based upon code performance and verified by experimental data from prototype laser systems. These specifications will optimize both performance and cost of future advanced laser systems.

IMAGING SOFTWARE

IMAGING SCIENCES WORKSHOP • 1994

IDP: Image and Data Processing (Software) in C++

Sean Lehman

IDP++ (Image and Data Processing in C++) is a compiled, multidimentional, multi-data type, signal processing environment written in C++. It is being developed within the Radar Ocean Imaging group and is intended as a partial replacement for View. IDP++ takes advantage of the latest object-oriented compiler technology to provide "information hiding." Users need only know C, not C++. Signals are treated like any other variable with a defined set of operators and functions in an intuitive manner. IDP++ is being designed for real-time environment where interpreted signal processing packages are less efficient.

A Probablistic Neural Network Classification System for Signal and Image Processing

Brenda Bowman

The Acoustical Heart Valve Analysis Package is a system for signal and image processing and classification. It is being developed in both Matlab and C, to provide an interactive, interpreted environment, and has been optimized for large scale matrix operations. It has been used successfully to classify acoustic signals from implanted prosthetic heart valves in human patients, and will be integrated into a commercial Heart Valve Screening Center. The system uses several standard signal processing algorithms, as well as supervised learning techniques using the probabilistic neural network (PNN). Although currently used for the acoustic heart valve application, the algorithms and modular design allow it to be used for other applications, as well. We will describe the signal classification system, and show results from a set of test valves.

An Object Oriented Environment for Multi-Channel Signal Analysis and Understanding

William J. Maurer, Farid U. Dowla

We describe an interactive signal analysis an understanding tool for multichannel signals. The system, written entirely in the C++ language, takes full advantage of the modern workstation GUI tools and integrates traditional signal-processing methods with intelligent domain-specific tools for the exploration and analysis of semistructured problems. By semistructured problems, we mean problems that require a high degree of interactive analysis, and further, the analysis steps are highly adaptive. In other words, a finite number of rules cannot be used to obtain a good solution to the problem.

The GUI interface allows rapid execution of powerful transformations and viewing of the transformation results. The system automatically stores the transformed data as objects for further inference and interpretation with a hierarchical knowledge-based representation. Hence, neural networks, fuzzy logic, or rule-based expert systems can be invoked on the knowledge-based objects. A persistent object capability allow users to save the knowledge-base objects in a file which can be loaded in the future or shared between users. Automatic parameterized signal- and image-processing algorithms are created by

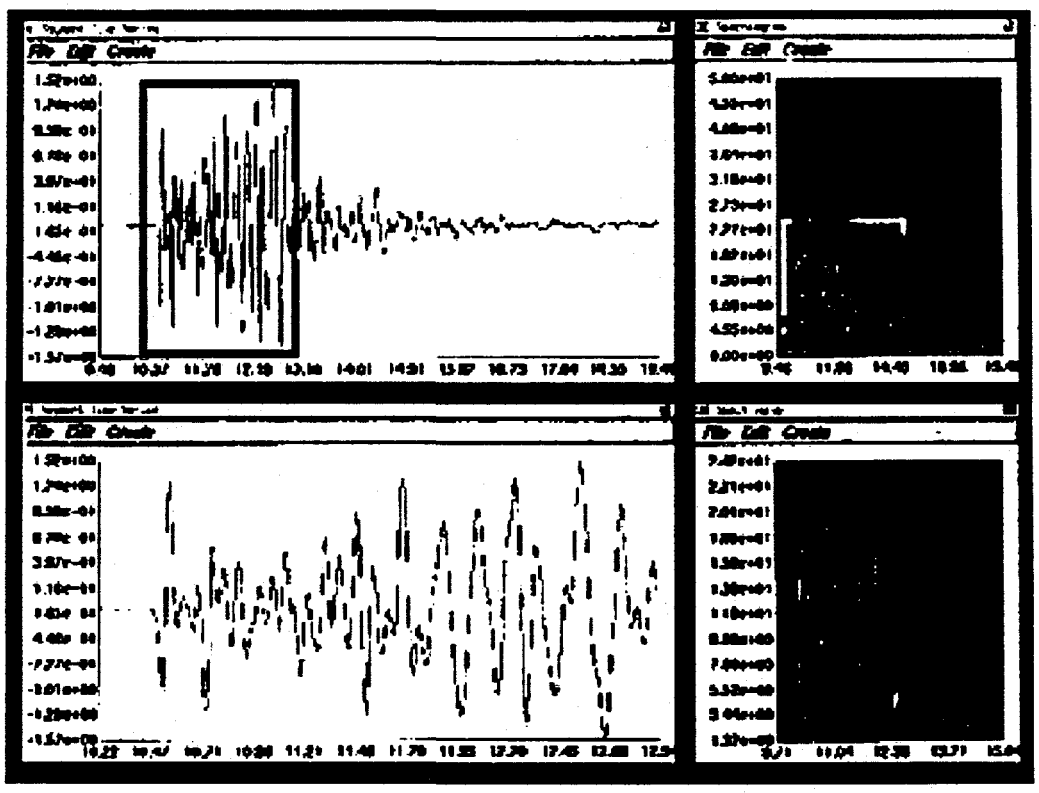

Figure 1. An example of interactive zooming in on a 1-D signal and a 2-D image. This spatial scaling capability is similar to the way a microscope allows you to focus in on an object of interest. Wavelet transforms will be used for time scaling capability. Using a combination of spatial and time scaling tools, features that are ambiguous may be studied carefully.

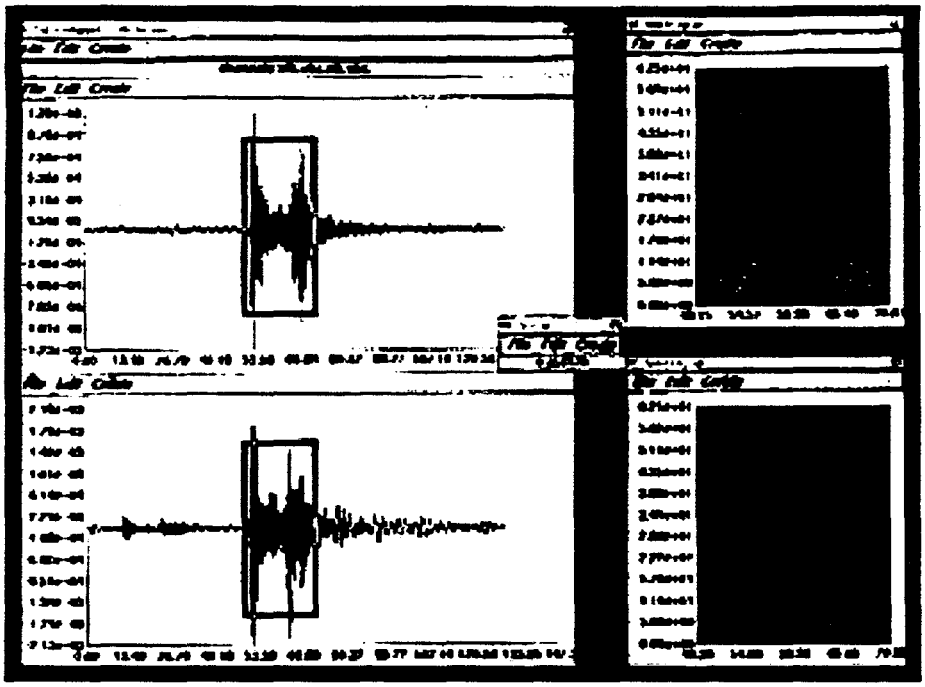

Figure 2. An example on interactive object oriented analysis of multichannel seismic signals. Using the system, the analyst is able to rapidly zoom in on sections of the data, correlation between the two spectrograms is computed resulting in the "Scalar" window indicating a value of the correlation coefficient.

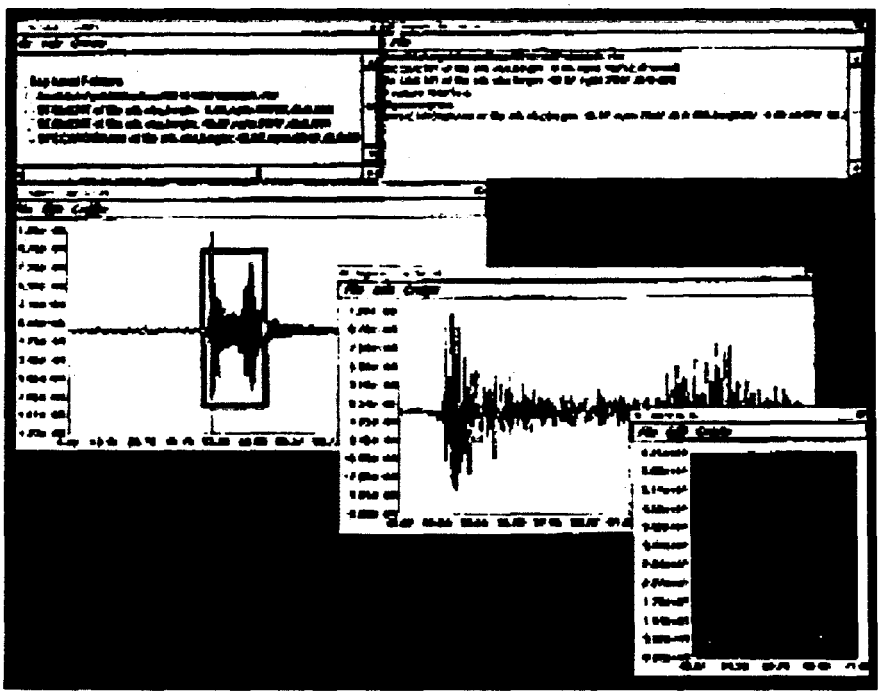

Figure 3. An example of deriving an object from the knowledge base. After importing data into a channel object, the user grows a multi-dimensional tree structure of selected segments and features. The tree, shown in the "Top Level Folders" window, is represented in a folder system similar to the Apple file system and allows the user to manage the objects created. A segment derivation, shown in the "Segment Derivation" window, is automatically stored and can be replayed at a later time.

ACOUSTIC / OCEANIC IMAGING

Current Issues in Radar Ocean Imaging

Dennis Holliday

A brief history of radar ocean imaging will be presented and current scientific issues in this subject will be summarized.

Radar Imagery from the 1994 Loch Linnhe Ship Wake Experiment

Carmen J. Mullenhoff, Sean K. Lehman, Holger Jones, Gary Berry, Tim Lamont-Smith, Phil Hirst, Kieth Ward

The 1994 Loch Linnhe radar ocean imaging trials were held from September 4 through September 17. Two ships were used: the R.V. Colonel Templer, and the RMAS Collie. Thorn EMI, Inc., fielded a dual band, dual polarization radar on a hillside overlooking the loch. A primary purpose of the experiment was to obtain highly visible images of ship generated internal waves. Presented here is imagery for a few of the good ship runs, as well as a study of the environment on the visibility of ship generated internal waves.

LLNL Current Meter Array—Concept and System Description

David D. Mantrom

Measurement Concept and Hardware Systems

A measurement capability using a horizontal array of 10 S4 current meters mounted on a stiff floating structure with 35 m aperture has been developed to support interpretation of radar imaging of surface effects associated with internal waves. This system has been fielded three times and most recently, has collected data alongside the sea-surface footprint of a land-fixed radar imaging ship-generated internal waves.

The underlying need for this measurement capability is described. The specifications resulting from this need are presented and the engineering design and deployment procedures of the platform and systems that resulted are described. The current meter data are multiplexed along with meteorological and system status data on board the floating platform and are telemetered to a shore station and on to a data acquisition system. The raw data are recorded, and are then processed to form space-time images of current and strain rate (a spatial derivative of the current field). Examples of raw and processed data associated with ship-generated internal waves are presented.

Near-surface Current Meter Array Measurements of Internal Gravity Waves

Holger B. E. Jones

Data Processing and Analysis

We have developed various processing algorithms used to estimate the wave forms produced by hydrodynamic Internal Waves. Furthermore, the estimated Internal Waves are used to calculate the Modulation Transfer Function (MTF) which relates the current and strain rate subsurface fields to surface scattering phenomenon imaged by radar. Following a brief discussion of LLNL's measurement platform (a 10 sensor current meter array) we describe the generation of representative current and strain rate space-time images from measured or simulated data. Then we present how our simulation capability highlighted limitations in estimating strain rate. These limitations spurred the application of beamforming techniques to enhance our estimates, albeit at the expense of collapsing our space-time images to 1-D estimates. Finally, we discuss progress with regard to processing the current meter array data captured during the recent Loch Linnhe field trials.

Determination of Radar MTF

David Chambers

The ultimate goal of the Current Meter Array (CMA) is to be able to compare the current patterns detected with the array with radar images of the water surface. The internal wave current patterns modulate the waves on the water surface giving a detectable modulation of the radar cross-section (RCS). The function relating the RCS modulations to the current patterns is the Modulation Transfer Function (MTF). By comparing radar images directly with co-located CMA measurements the MTF can be determined.

In this talk radar images and CMA measurements from a recent experiment at Loch Linnhe, Scotland, will be used to make the first direct determination of MTF for an X and S band radar at low grazing angles. The technical problems associated with comparing radar images to CMA data will be explained and the solution method discussed. The results suggest the both current and strain rate contribute equally to the radar modulation for X band. For S band, the strain rate contributes more than the current. The magnitude of the MTF and the RCS modulations are consistent with previous estimates when the wind is blowing perpendicular to the radar look direction.

Laser Sources for Object Illumination

Georg F. Albrecht

Abstract

The considerations which formulate the specifications for a laser illuminator are explained, using the example of an underwater object. Depending on the parameters which define the scenario, widely varying laser requirements result.

Laser Sources for Object Illumination

Numerous applications require the illumination of the object to be imaged by an external source. To the extend that the illumination is artificial (as opposed to sun light), it becomes one more parameter to be optimized with respect to the envisioned application. Among the possibilities are timed short pulses to selectively illuminate and then image objects at a specified distance, illumination at one or several selected wavelengths to enhance contrasts in the image or enable propagation of the illumination within a spectral window, illumination which induces species specific fluorescence in part of the object to be imaged, or one can determine its shape by illuminating it with a fringe pattern generated by two lasers of lightly different narrow band frequencies. In what follows the thought process determining the specifications of a typical illumination systems and the technology involved will be briefly described.

Given that illumination over appreciable distances always translates into a requirement for source radiance (W/cm^2 sterad), a laser is mostly the illuminator of choice. If a specific wavelength is an additional requirement, lasers quickly become the only possibility. By the time one adds the need for compactness and hands-off reliability like one typically encounters in systems on pods and remotely operated units the options have further narrowed to solid state lasers. To appreciate some of the demands put on illuminators, let us consider a system which detects objects under water.

Figure 1 shows the transparency of different kinds of seawater. Most obviously, a wavelength of around .5 μm has the lowest propagation losses. Imaging detectors are well developed in this wavelength region.

Figure 2 shows a typical scenario. A transceiver system resides above the water, attempting to locate an underwater object and determine its shape and depth. Information on the outline will be contained in an image of reflected illuminator light, depth information will be contained by the temporal arrival of the illuminator return from the object of interest. CW illumination obviously contains no depth information. As indicated in the figure, there are a variety of other effects returning reflected light, mainly the water surface, and bulk scattering within the water column. Hence the return signal for an opaque reflector generically looks like in the insert to Figure 2, where the first return corresponds to the water surface, the tail corresponds to scattering from different depths along the water column as the illuminator light travels the object, and the last feature corresponds to the object itself.

Now one can see some of the demands such a scenario puts on the illuminator. Firstly, it needs to have a wavelength of around .5 μm. Secondly, the pulse duration not only determines the depth resolution, but also determines the minimum distance an object must lie under the water surface so that a return from it can be gated out from the water surface return itself, and thus its depth be determined. Thirdly, the laser pulse energy together with the detector sensitivity and background effects determines down to what depth an object can be "seen" by illuminating it. Finally, the pulse repetition rate is determined by the desired search rate. If the illuminator application is over large distances through the atmosphere or through space, the beam quality of the laser output becomes one of the most crucial parameters of all.

In practice, the desired wavelength is best obtained by frequency doubling a Nd^{3+} based solid state laser to ≈ .53 μm. An acceptable pulse duration (a few tens of ns) is achieved by a technique called Q-switching, for even shorter pulses (few ns) cavity dumping can be employed, although its technical demands on the laser cavity switching element are considerably more refined. A sufficient laser pulse energy can in principle be obtained by amplification, but the packaging demands (compactness) associated with autonomous systems quickly turn this into a rather sophisticated engineering problem. The pulse repetition rate surfaces in the laser as a thermal management problem, and in the overall system as a prime power issue which quickly turns into demands of maximized efficiencies wherever possible.

As much as the above example applies for an under water scenario, it illustrates the type of considerations which determine the specification of the illuminating laser source. Note that since the laser delivers the light ultimately received by the detector, it is not simply a stand-alone item, but in any application laser power, beam

quality and wavelength can be traded off against signal to noise, intrinsic detector sensitivity, background suppression methods at the receiver end, receiver aperture and field of view.

Compared to other laser applications, illuminators with ranging capability sometimes demand rather sophisticated pulse formats and reasonably large powers. By the time the system is engineered to compactness and exhibits a degree of reliability which keeps it out of the user's mind, such systems often come to represent quite the state of the art. Applications such as the one described above have, in fact, helped to push laser technology in the past and it is to be hoped that they will in the future.

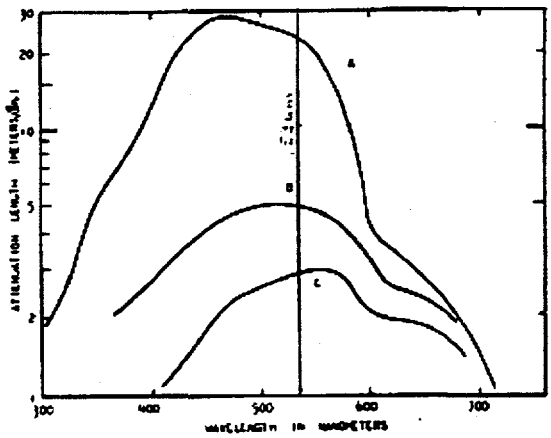

Figure 1. Attenuation length in three types of sea water (from CRC Handbook of Lasers, CRC Press, 1971). Also indicated in the wavelength of .53 μm of frequency doubled Nd^{3+} solid state lasers.

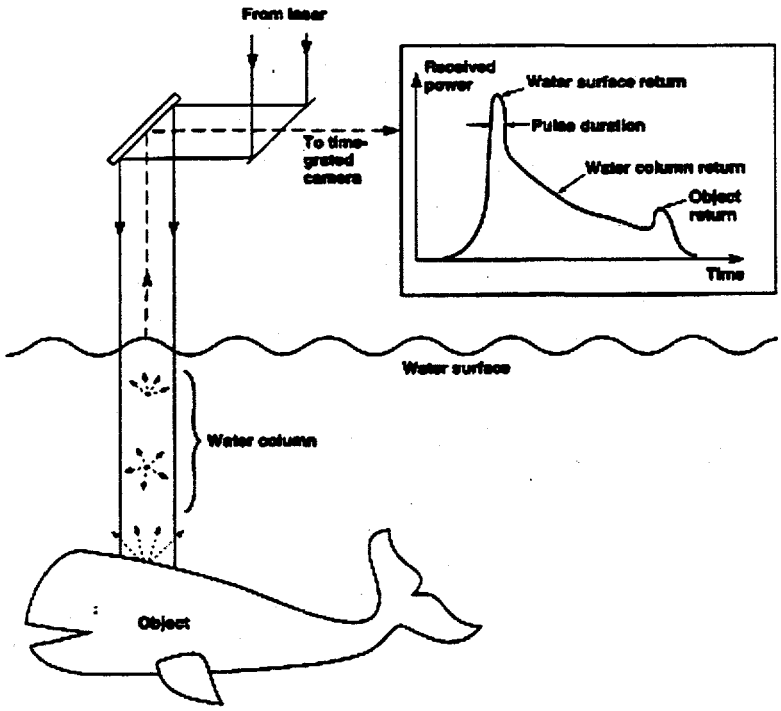

Figure 2. A typical scenario generates returns from the water surface, the water column between the surface and the object, and the object. This is reflected in the temporal shape of the return as shown in the insert.

Underwater Laser Imaging System (UWLIS)

Michael DeLong

Practical limitations with underwater imaging systems are reached when the noise in the back scattered radiation generated in the water between the imaging system and the target obscures the spatial contrast and resolution necessary for target discovery and identification. The advent of high power lasers operating in the blue-green portion of the visible spectrum (oceanic transmission window) has led to improved experimental illumination systems for underwater imaging. Range-gated and synchronously scanned devices take advantage of the unique temporal and spatial coherence properties of laser radiation, respectively, to overcome the deleterious effects of common volume back scatter.

Synchronously scanned systems rely on the highly collimated nature of the laser beam for spatial rejection of common-volume back scatter. At Lawrence Livermore National Laboratory we have developed a synchronous, raster-scanning underwater laser imaging system (UWLIS). The present UWLIS system differs from earlier synchronous scanners in its ability to scan in two dimensions at conventional video frame rate (30 Hz). The performance of the present UWLIS during an in-water tank test is summarized in which the imaging performance of the system was tested at distances of up to 6.3 AL (at a physical distance of 15.2 m). The test results indicate that the UWLIS system already is capable of extending the underwater imaging range beyond the capabilities of conventional floodlight-illuminated SIT camera systems. The ultimate range achievable for this type of system is determined by the laser power, optical collection geometry, and the degree of frame averaging that is used.

The two-dimensional imaging capabilities of UWLIS fill specific need in underwater salvage and mine detection operations that is not possible with other synchronous-scanning systems, such as line scanners. The real or near-real time frame rates of the UWLIS make possible operation in a stationary mode or in a mode in which the platform speed is randomly varied. This is typical of the operational environment, in which the platform is often maneuvered above and around rugged seafloor terrains. Signal-to-noise and resolution requirements do, however limit the system to relatively narrow fields-of-view. In this respect the UWLIS can be considered a complementary device to line scanners, which permit larger swath scans but require forward vehicle motion and lower horizontal line rates.

Structural Acoustic Imaging Techniques

Terry Donich

Abstract not available.

Multispectral Image Feature Fusion for Detecting Land Mines

Gregory A. Clark, Sailes K. Sengupta, William D. Aimonetti, Frank Roeske, John G. Donetti, David J. Fields, Robert J. Sherwood, Paul C. Schaich

Our system fuses information contained in registered images from multiple sensors to reduce the effects of clutter and improve the ability to detect surface and buried land mines. The sensor suite currently consists if a camera that acquires images in sixible wavelength bands, du, dual-band infrared (5 micron and 10 micron) and ground penetrating radar. Past research has shown that it is extremely difficult to distinguish land mines from background clutter in images obtained from a single sensor. It is hypothesized, however, that information fused from a suite of various sensors is likely to provide better detection reliability, because the suite of sensors detects a variety of physical properties that are more separable in feature space. The materials surrounding the mines can include natural materials (soil, rocks, foliage, water, holes made by animals and natural processes, etc.) and some artifacts.

We use a supervised learning pattern recognition approach to detecting the metal and plastic land mines. The overall process consists of four main parts: preprocessing, feature extraction, feature selection, and classification. These parts are used in a two step process to classify a sub-image. The ***first step***, referred to as feature selection, determines the features of sub-images which result in the greatest separability between the classes. The ***second step***, image labeling, uses the selected features and the decisions from a pattern classifier to label the regions in the image which are likely to correspond to mines.

We extract features from the images, and use feature selection algorithms to select only the most important features according to their contribution to correct detections. This allows us to save computational complexity and determine which of the spectral bands add value to the detection system. The most important features from the various sensors are fused using a supervised learning pattern classifier (the probabilistic neural network). We present results of experiments to detect land mines from real data collected from an airborne platform, and evaluate the usefulness of fusing feature information from multiple spectral bands. We show that even with preliminary data and limited testing, the performance (specified in terms of probability of detection and probability of false alarm) is very promising. The novelty of the work lies mostly in the combination of the algorithms and their application to the very important and currently unsolved operational problem of detecting minefields from an airborne standoff platform.

"Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract number W-7405-ENG-48."

LIFTERS-Hyperspectral Imaging at LLNL

David Fields, Charles Bennett, Mike Carter

LIFTIRS, the Livermore Imaging Fourier Transform InfraRed Spectrometer, recently developed at LLNL, is an instrument which enables extremely efficient collection and analysis of hyperspectral imaging data. LIFTIRS produces a spatial format of 128x128 pixels, with spectral resolution arbitrarily variable up to a maximum of 0.25 inverse centimeters. Time resolution and spectral resolution can be traded off for each other with great flexibility. We will discuss recent measurements made with this instrument, and present typical images and spectra.

MICROWAVE / ACOUSTIC IMAGING

IMAGING SCIENCES WORKSHOP • 1994

Micropower Radar Systems for Law Enforcement Technology

Stephen G. Azevedo, Jeff Mast, Mel Wieting, T. McEwan, J. Trebes, J. Brase

LLNL researchers have pioneered a unique compact low-power and inexpensive radar technology (microradar) that has enormous potential in various industries. Some licenses are currently in place for motion sensors and stud finders. The ultra-wideband characteristics of the microradar (4 to 10 GHz) make it difficult to detect, yet provide potential range resolution of 1 cm at ranges of greater than 20 meters. Real and synthetic apertures arrays of radar elements can address imaging applications behind walls at those distances. Personnel detection applications are currently being tested.

Three-Dimensional Ground Penetrating Radar Imaging Using Multi-Frequency Diffraction Tomography

Jeffrey E. Mast, Erik M. Johansson

In this talk we present results from a three-dimensional image reconstruction algorithm for impulse radar operating in monostatic pule-echo mode. The application of interest to us is the nondestructive evaluation of civil structures such as bridge decks. We use a multi-frequency diffraction tomography imaging technique in which coherent backward propagations of the received reflected wavefield form a spatial image of the scattering interfaces within the region of interest. This imaging technique provides high-resolution range and azimuthal visualization of the subsurface region. We incorporate the ability to image in planarly layered conductive media and apply the algorithm to experimental data from an offset radar system in which the radar antenna is not directly coupled to the surface of the region. We present a rendering in three-dimensions of the resulting image data which provides high-detail visualization.

Finite-Difference Modeling of Commercial Aircraft Using TSAR

Steven T. Pennock, Andrew J. Poggio

Future aircraft may have systems controlled by fiber optic cables, to reduce susceptibility to electromagnetic interference. However, the digital systems associated with the fiber optic network could still experience upset due to powerful radio stations, radars, and other electromagnetic sources, with potentially serious consequences.

We are modeling the electromagnetic behavior of commercial transport aircraft in support of the NASA Fly-by-Light/Power-by-Wire program, using the TSAR finite-difference time-domain code initially developed for the military. By comparing results obtained from TSAR with data taken on a Boeing 757 at the Air Force Phillips Lab, we hope to show that FDTD codes can serve as an important tool in the design and certification of U.S. commercial aircraft, helping American companies to produce safe, reliable air transportation.

Visualization of Elastic Wavefields Computed with a Finite Difference Code

Shawn Larsen, Dave Harris

Intelligence and National Security Technologies Program

We have developed a finite difference elastic propagation model to simulate seismic wave propagation through geophysically complex regions. To facilitate debugging and to assist seismologists in interpreting the seismograms generated by the code, we have developed an X Windows interface that permits us to view successive temporal snapshots of the (2D) wavefield as they are calculated. A sequence of snapshots can be stored to a file, then replayed as a "movie". To assist interpretation further, the compression and shear potentials are separated and displayed, superimposed, in different colors. To simplify rendering the superimposed potentials on 8-bit displays, the potentials are rendered separately on offset grids following the pattern:

PSPSPSPSPSPSP

SPSPSPSPSPSPS

PSPSPSPSPSPSP

SPSPSPSPSPSPS

We present a brief video displaying the generation of seismic waves by an explosive source on a continent, which propagate to the edge of the continent then convert to two types of acoustic waves. This sample calculation was part of an effort to study the potential of offshore hydrocoustic systems to monitor seismic events occurring onshore.

The elastic propagation model is a highly optimized solver for the 2D elastodynamic formulation of the full wave equation on a staggered grid. This code is 4th order accurate in space and 2nd order accurate in time. Absorbing boundary conditions using low-order paraxial methods are applied to the side and bottom grid interfaces, which effectively reduce artificial boundary reflections by 90-95%. The algorithm vectorizes well with a performance of over 300 MFlops on a single processor of a Cray Y-MP. A previous version of the code has been parallelized to run on a 192 node nCUBE-2 parallel processor, made available by the Earth Resources Laboratory at the Massachusetts Institute of Technology (up to 128 nodes can be used simultaneously). The code also features a variable density grid option. This is particularly useful in seismic propagation problems, where very-low-velocity surface layers (e.g. water) overlie higher-velocity deep structures. A relatively dense grid is required in the shallow layers to prevent numerical dispersion; a coarse grid suffices in the high-velocity deep layers. The code presently allows three gradations in gridding density resulting in considerable savings in computer resources. We also incorporate an active-grid algorithm which eliminates calculations in regions void of seismic energy. For typical problems, the variable-density grid and active-grid algorithms can reduce memory requirements by a factor of 10 and run times by a factor of 20 above and beyond the vectorization and parallelization performance gains.

Current enhancements to the code include a port to the Lab's 256-processor Meiko CS-2, the incorporation of viscoelastic attenuation, an axi-symmetric 3D adaptation (cylindrical coordinates) and a full 3D implementation. The single-node Meiko performance on a similar acoustic code is better than 150 MFlops; we have achieved over 26 GFlops with the acoustic code using 252 nodes of the machine. Comparable speeds are expected for the elastic code. The code is being enhanced to incorporate attenuation through the addition of memory variables.

The viscoelastic and 3D versions of the code will be applied to Treaty Verification modeling tasks and to seismic risk assessment. For the verification application, the code will model wave propagation in strongly heterogeneous (tectonically active) regions of interest such as the Middle East and North Africa, and to examine the use of seismic systems on islands to detect waterborne T phases (long range observations of explosions at sea). For seismic risk assessment, the code will simulate focusing of seismic energy in low-velocity sediment-filled basins such as the Marina District of San Francisco. It is important to determine the geographical distribution of such high-amplitude seismic ground motion. We propose to investigate the effect of 3-D crustal structure on seismic wave propagation, generalizing on most previous research in this field which has been limited to 2-D numerical simulations.

COMPUTED TOMOGRAPHY

IMAGING SCIENCES WORKSHOP • 1994

A Constrained Conjugate Gradient Algorithm for Computed Tomography

Stephen G. Azevedo, Dennis M. Goodman

Image reconstruction from projections of x-ray, gamma-ray, protons and other penetrating radiation is a well-known problem in a variety of fields, and is commonly referred to as computed tomography (CT). Various analytical and series expansion methods of reconstruction and been used in the past to provide three-dimensional (3D) views of some interior quantity. The difficulties of these approaches lie in the cases where (a) the number of views attainable is limited, (b) the Poisson (or other) uncertainties are significant, (c) quantifiable knowledge of the object is available, but not implementable, or (d) other limitations of the data exist.

We have adapted a novel nonlinear optimization procedure developed at LLNL to address limited-data image reconstruction problems. The technique, known as nonlinear least squares with general constraints or constrained conjugate gradients (CCG), has been successfully applied to a number of signal and image processing problems, and is now of great interest to the image reconstruction community. Previous applications of this algorithm to deconvolution problems and x-ray diffraction images for crystallography have shown the great promise.

The non-linear, constraint-based conjugate gradient optimization method is an iterative algorithm, so it can be relatively slow compared to other reconstruction methods. Our implementation has been optimized to reduce the reconstruction time, and its improved convergence properties over unconstrained methods make competitive in overall speed. We will describe the CCG algorithm and show results of its use in several limited-data CT problems for use in nondestructive evaluation. New features to be described are the use of Poisson noise statistics, adaptation of the inverse problem, and extensive incorporation of known constraints.

Ultra High Resolution Tomography

W. S. Haddad

Recent work and results on ultra high resolution three dimensional imaging with soft x-rays will be presented. This work is aimed at determining microscopic three dimensional structure of biological and material specimens. Three dimensional reconstructed images of a microscopic test object will be presented; the reconstruction has a resolution on the order of 1000 Å in all three dimensions. Preliminary work with biological samples will also be shown, and the experimental and numerical methods used will be discussed.

Applications of Electrical Resistance Tomography To Subsurface Environmental Restoration

A.L. Ramirez, W.D. Daily

We are developing a new imaging technique, Electrical Resistance Tomography (ERT), to map subsurface liquids as flow occurs during natural or clean-up processes and to map geologic structure. Natural processes (such as surface water infiltrating the vadose zone) and man-induced processes (such as tank leaks and clean-up processes such as steam injection, can create changes in a soil's electrical properties that are readily measured. We have conducted laboratory and a variety of field experiments to investigate the capabilities and limitations of ERT for imaging underground structures and processes. In the last four years we have used ERT to successfully monitor several field processes including: a subsurface steam injection process (for VOC removal), an air injection process (below the water table) for VOC removal, water infiltration through the vadose zone, radio-frequency heating, ohmic heating, and tank and pond leaks. The information derived from ERT can be used by remediation projects to: detect and locate leaks, determine the effectiveness of clean-up processes, select appropriate clean-up alternatives, and to verify the installation and performance of subsurface barriers.

Description of ERT

To image the resistivity distribution between two boreholes, we placed a number of electrodes in electrical contact with the soil in each borehole Using an automatic data collection and switching system we then applied a known current to any two electrodes and measured the resulting voltage difference between other pairs of electrodes. Each ratio of measured voltage and current is a transfer resistance. Next, we switched to two other electrodes, applied current between two other electrodes and again measured the voltage differences using electrode pairs not being used for the source current. We repeated this process until many combinations were measured which completely encircled the target area. For n electrodes there are $n\,(n\,-3)/2$ linearly independent transfer resistances. A complete set of linearly independent data contains the maximum information content about the target; any additional measurements collected are redundant.

Four point measurements were used to eliminate the effect of electrode contact resistance on the measured

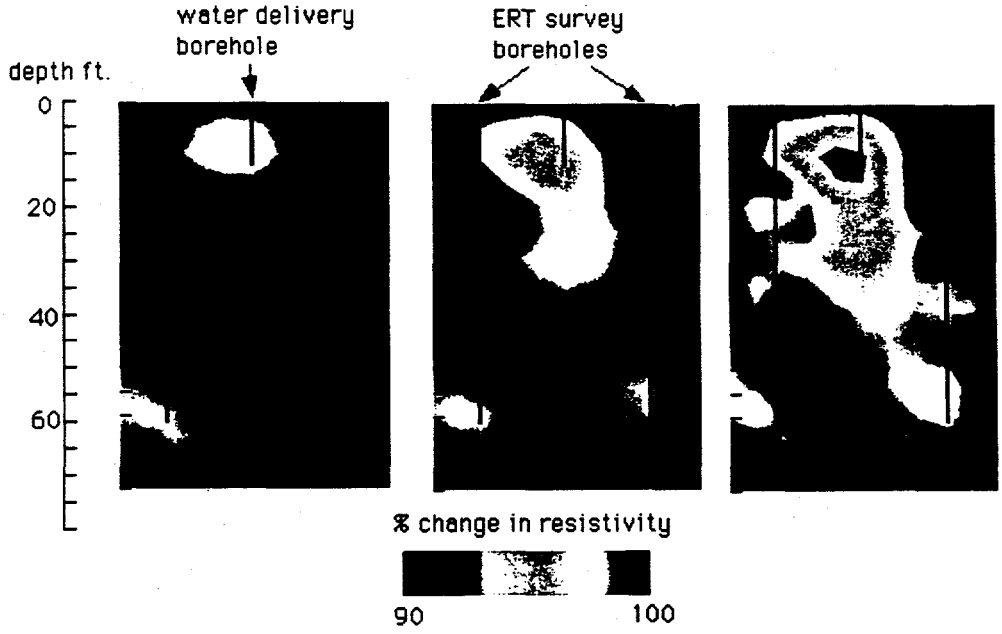

Figure 1. These ERT tomographs show changes in soil resistivity caused by water as it moves along subsurface flow paths in the unsaturated zone.

values of formation resistance. The measurements were made using a direct current (DC) measurement system which switched the polarity of the transmitter voltage and corrected for naturally occurring self-potentials. The DC frequency eliminates inductive coupling between wires connecting each electrode in a borehole to the ground surface. The transmitter-receiver combinations sampled provided a complete set of linearly independent measurements as well as some redundant measurements. The combinations sampled included transmitter-receiver pairs located within the same borehole, transmitters in one borehole and receivers in the other borehole, and combinations which had each pole of the receiver and of the transmitter pair located in separate boreholes.

Calculating the distribution of resistivity in the vicinity of the boreholes given the measured transfer resistances is a highly nonlinear problem because the current paths are dependent on the resistivity distribution. This type of inversion has been widely studied by others. We use a finite element approach to solve the forward problem. The inverse problem is then solved as an optimization problem where an objective function is minimized subject to adequately fitting the measurements.

Some of the desirable traits of ERT as a new environmental technology are as follows: 1-The technique minimizes the need for drilling monitoring wells while providing unprecedented detail of subsurface structures and processes. 2-Data acquisition rate is sufficient to allow detailed sampling of large subsurface regions in reasonable amounts of time. 3-Sensors are cheap and robust—pieces of stainless steel buried in the ground to depths of hundreds of feet; sensors can be placed in boreholes (alternatively they can be pushed into the ground thereby reducing need for drilling) and/or at the ground surface. 4- Borehole separation can range from a few tens to hundreds of feet. 5-The technology can be applied to a wide range of conditions and sites; it has been demonstrated in a variety of geologic environments.

PHYSICAL IMAGING

IMAGING SCIENCES WORKSHOP • 1994

Quantifying the Topology of Porous Structures

John Kinney

Abstract

Computerized x-ray tomography, with microscopic resolution, has been used to volumetrically visualize the evolution of porosity in a ceramic matrix composite during processing. The topological variables describing the porosity have been measured. The evolution of the porosity exhibits critical scaling behavior near final consolidation, and appears to be independent of the structure (universality).

Imaging Surface Atomic Structure with Energy Dependent Electron Diffraction

Jim Tobin

Abstract

The general concepts and approaches to extracting structural information from imaging with electron diffraction are outlined.

Processing Multidimensional Nuclear Physics Data

John Becker

Modern Ge detector arrays for gamma-ray spectroscopy are producing data sets unprecedented in size and event multiplicity.

Gammasphere, the DOE sponsored array, has the following characteristics:

- High granularity (110 detectors)
- High efficiency (10%)
- Precision energy measurements (Delta EE = 0.2%)

Characteristics of detector line shape, the data set, and the standard practise in the nuclear physics community to the nuclear gamma-ray cascades from the 4096 times 4096 times 4096 data cube will be discussed.

Determination of Chemical Concentration with a 2 Dimensional CCD Array in the Echelle Grating Spectrometer

D. Kent Lewis, Charles G. Stevens

Abstract

The Echelle grating spectrometer (EGS) uses a stepped Echelle grating, prisms and a folded light path to miniaturize an infrared spectrometer. Light enters the system through a slit and is spread out along Y by a prism. This light then strikes the grating and is diffracted out along X. This spreading results in a superposition of spectral orders since the grating has a high spectral range. These orders are then separated by again passing through a prism. The end result of a measurement is a 2 dimensional image which contains the folded spectrum of the region under investigation. The data lies in bands from top to bottom, for example, with wavenumber increments as small as 0.1 lying from left to right such that the right end of band N is the same as the left end of band N+1. This is the image which must be analyzed.

The analysis process is then:

(1) Acquire an image

(2) Identify the wavenumbers contained in each region of the image

(3) Build a model for the image from a library function

(4) Compare the model to the measurement

(5) Deconvolve the measurement image segment by segment

(6) Determine the composition and concentration of the enhanced spectrum.

The system effects which are included in the measurement are:

(1) Pixelization, distortion caused by measuring a continuous function over a fixed length by summation over the pixel area,

(2) Slit diffraction, which causes distortion of the idealized slit image,

(3) Optical aberration, the misalignment of data caused by realizable optical systems, and

(4) Atmospheric differences from the idealized case caused by unforeseen atmospheric conditions. Notice that the atmosphere itself is considered part of the measurement system.

The analytical processor is model based in that an idealized measurement is created from library spectra. This model will be compared to the measured spectral image. The measurement is separated into each widened spectral line and a Gaussian fit to the cross spectral line shape to smooth out the system irregularities.

This also allows a single measurement line to be extracted from each region by averaging the smoothed data.

In any measurement, there are atmospheric lines which will not be complicated by the molecules which are being sought. These lines are compared between the measurement and library function to determine the state of the atmosphere itself. In case the measurement is very far from the expected values, the process of significant factor analysis can be used to determine the number of chemical species present, perhaps showing the need for a wider library of reference functions.

This step identifies line widening and water attenuation caused by heat of moisture in the air. The library function can then be changed to more closely match the actual conditions.

The reference spectrum is broken up and overlaid in a 2 dimensional matrix. An idealized slit function is then cross multiplied by these spectral lines. This image is the model to which the measured image is compared.

The two images are then broken into overlapping segments and deconvolved using a 2 dimensional Weiner filter. The slit diffraction function and pixelization function are known in general, so this part of the filter may be pre-calculated and simply checked in the algorithm. This process accentuates the differences between the two images.

We have now corrected for irregularities in the slit transfer function, the pixelization smearing, the atmospheric broadening, and any optical aberration present. We have done this region by region, and these regions can be made variable in case the analysis later reveals too high an uncertainty in chemical species determination.

In case the optical aberration is too great and is different in different ranges of the image, a state space deconvolver can be employed to model a signal (the transmission spectrum line) convolved by a system which changes in time. This would allow a better solution in general, but its overhead costs in terms of time would seem to limit this technique to extreme cases.

The data is now in a form which we can analyze, and the various techniques of chemometrics have been explored to find an optimal method. All those considered use, to some extent, an eigenvalue / eigenvector solution and for the very large, sparse data sets that the EGS typically generates, we must use a singular value decomposition method to handle the problem.

Another possible modification which could simplify the solutions is remeshing the data into a more dense, more square matrix. This has the disadvantage of requiring exact matching of the ends of the as yet unknown data.

When properly modified, all the techniques work well. Those investigated have been linear least squares, partial least squares, nonlinear least squares, K matrix and P matrix solvers.

Infrared Source Test

Linda Ott

The purpose of the Infrared Source Test (IRST) is to demonstrate the ability to track a ground target with an infrared sensor from an airplane. The system is being developed within the Advance Technology Program's Theater Missile Defense/ Unmanned Aerial Vehicle (UAV) section.

The IRST payload consists of an Amber Radiance 1 infrared camera system, a computer, a gimbaled mirror, and a hard disk. The processor is a custom R3000 CPU board made by Risq Modular Systems, Inc. for LLNL.

The board has ethernet, SCSI, parallel I/O, and serial ports, a DMA channel, a video (frame buffer) interface, and eight MBytes of main memory. The real-time operating system VxWorks has been ported to the processor. The application code is written in C on a host SUN 4 UNIX workstation. The IRST is the result of a combined effort by physicists electrical and mechanical engineers, and computer scientists.

The closed-loop tracking cycle consists of capturing the digital image, locating the target in the image, adjusting the image exposure control, calculating the pointing vectors, and commanding the gimbal azimuth and elevation positions. The current cycle time for this loop is fifty (50) milliseconds. The image array is 256 by 256 pixels. The results of tracking a target in the laboratory will be presented. During the flight experiment, the digital images will be stored on a local disk; the analog video and tracking parameters will be telemetered to the ground for real-time display. This system is another link in the evolution of both hardware and software developed for other ATP programs: ARGUS, Brilliant Pebble Flight Experiments, sensor calibration, and Clementine

IMAGING ALGORITHMS

IMAGING SCIENCES WORKSHOP • 1994

Comparison of Image Deconvolution Algorithms on Simulated and Laboratory Infrared Images

Deanne Proctor

Abstract

We compare Maximum Likelihood, Maximum Entropy, Accelerated Lucy-Richardson, Weighted Goodness of Fit, and Pixon reconstructions of simple scenes as a function of signal-to-noise ratio for simulated images with randomly generated noise. Reconstruction results of infrared images taken with the TAISIR (Temperature and Imaging System InfraRed) are also discussed.

Summary

Software has been developed to compare image reconstruction results for Maximum Likelihood, Maximum Entropy, Accelerated Lucy-Richardson, Goodness of Fit, and Pixon methods of image reconstruction. For a given truth image, point spread function, and noise standard deviation, the procedure convolves the truth image with the point spread function and adds randomly generated noise with Gaussian distribution of input standard deviation. The resulting simulated image is then used as input to the above deconvolution procedures and the mean square error for (reconstruction result - truth image) is calculated as well the maximum error of same. Summary plots of an arbitrary slice of each reconstruction are generated.

This software was used to examine simple scenes with various values of signal-to-noise ratio and various object dimensions. A well-sampled apodized-aperture point spread function obtained for the TAISIR flight system was used, with seven pixels across the central peak. No noise was added to the point spread function. To date, object sizes of 10, 30 and 50 pixels have been studied. Defining signal as the difference between object and background, signal/noise ratios of 1000, 100, 30, and 10 were examined.

For signal-to-noise ratio of 1000 all methods exhibit convergence. At sigma-to-noise ratio of 100, some spurious signal is noticed with all methods but Pixon. By signal-to-noise ratio of 30, strong loss of fidelity is apparent with Maximum Likelihood and Maximum Entropy. At signal-to-noise ratio of 10, Accelerated Lucy-Richardson has also begun to "grow" noise into spurious signal, though the iterations are required to reconstruct the true signal to its appropriate value. Overall, for the types of scenes examined, the Pixon method resulted in reconstructions that best maintained fidelity to the truth image.

**All software was written in IDL (Interactive Data Language), a product of Research systems, Inc.*

**Maximum Entropy and Maximum Likelihood procedures were written by Frank Varosi*

STX @ NASA/GSFC and distributed with the IDL Astronomy Users Library.

Questions about the IDL Astronomy Library can be addressed to:

Wayne Landsman, Internet address: landsman@stars.gsfc.nasa.gov.

**Goodness of Fit and Pixon procedures were developed by R.C. Puetter and R.K. Pina. See Pixons and Bayesian Image Reconstruction, Proc. SPIE, Jul. 94.*

Internet address: rpuetter@ucsd.edu

**The Accelerated Lucy-Richardson routine is copyrighted by Hans-Martin Adorf & Fionn Murtagh. Internet address: adorf@eso.org and fmurtagh@eso.org.*

See "Accelerating the Richardson-Lucy restoration algorithm", Proc. 4th ESO/ST-ECF Data Analysis Workshop, 13-15 May 1992, Garching b. Muenchen, P. Grosboel et al. (eds.)

Some Topics in the Spatial Bispectra

Edmund Sullivan

The bispectrum can be defined as the triple fourier transform of the third order cumulant of a data series. Up to the present, except in image analysis, most work on the bispectrum has treated time series. Recently however, there has been interest in using the bispectrum in acoustic array processing. After a look at some issues involving sampling frequencies and symmetries of the bispectrum in general, two applications of the spatial bispectrum to underwater acoustic array processing will be discussed. One is a method of processing against loss of spatial coherence in towed arrays, which takes the form of a one-dimensional image, and the other is a look at the role of spatial bispectra in matched-field processing, which is a form of model-based processing used for the localization of acoustic sound sources.

Use of the Wavelet Transform in Signal Interpretation

Farid Dowla

Abstract not available.

Integer Anatomy

R. Doolittle

I. Introduction and Preliminary Form

The title integer anatomy is intended to convey the idea of a systematic method for displaying the prime decomposition of the integers. Just as the biological study of anatomy does not teach us all things about behavior of species neither would we expect to learn everything about the number theory from a study of its anatomy. But, some number-theoretic theorems are illustrated by inspection of integer anatomy, which tend to validate the underlying structure and the form as developed and displayed in this treatise. The first statement to be made in this development is: the way structure of the natural numbers is displayed depends upon the allowed operations.

In experimental physics, a field must be applied to a quantity before its structure can be determined. Thus it could be argued that the structure of the quantity is conditioned upon the applied field and that our perception of reality does not exist for either, independently. For it is just as true that a field lacks a description unless there is a particle or other quantity for it to act on. A famous example is the revealing of DNA linear structure through the application of fields at right angles which unwrapped the helix that bound them.

Thus we associate structure in numbers with a set of allowed operations. The operation of addition produces the "natural" ordering of the integers, which are conveniently arrayed for display as a sequence that proceeds left to right from a starting number. The prescription (we use the term "prescription" rather than "algorithm" since the instructions given are not necessarily complete in all detail to accomplish the task) for producing the so-called counting sequence of integers would be:

1) Start with zero

2) Add 1, algebraically, to the right-most number

3) Append the result to the right

4) Go to step 2.

We could call this display the counting form or "c-form" of the integers. We compare this form with any other form to establish a one-to-one correspondence with the integers. This is accomplished, for instance, by lining up a test sequence in say a row that is below the c-form row. There are other forms akin to this form, revealed by allowing other operations, as we will see.

Adding the operation of subtraction introduces the negative integers which are displayed as a lineup of integers proceeding to the left from zero. In the above prescription "add 1" is replaced by "subtract one" while "left" replaces "right". The operation of addition is now valid over the whole combined sequence as is subtraction. Until the idea of subtraction was introduced, the integers needed no modifier such as "positive" or "negative". Thus the structure that describes the integers has taken on more complexity as the number of allowed operations is increased, suggesting a link between the operations and description of structure.

Further complexities are introduced by the operation of multiplication and division. It is these two operations, in particular, that help produce the operation-dependent description we wish to explore further.

The operation of division, when constrained to yield an integer, reveals a highly significant limitation—it is applied to pairs of integers; i.e., it is "2-ary" but not all pairs of integers will produce another integer. Thus the integers fail to be holonomic under this constraint. This has had many consequences and is a major concern of number theory. My concern is limited, for now, to the set of prime numbers. These are numbers which have only unity or themselves as "proper" divisors; i.e., those divisors that yield an integer.

While it took the operation of division to recognize the existence of prime numbers, it now takes the operation of multiplication to reveal another conditional structure of the integers. The structure of the integers conditioned upon multiplication is expressed by the fact that any integer has its equal in some product of prime numbers, each prime being raised to some integer power. This is a well-known theorem whose proof is not needed here. Call this the prime-structure of the integers or simply "p-form". Just as with the additive structure mentioned above, or "c-form" it is convenient to find a way to display the p-form of the integers. It is conjectured that a very convenient display of the p-form exists.

Before this form is shown. it is instructive to consider an intermediate and very recognizable form, call it the "f-form". This form is displayed by a very simple prescription as follows:

1) Display two zeros

0 0

2) Add 1 algebraically to the right-most number and replace it with the result.

0 1

3) Copy once and append to the right,

0 1 0 1

4) Repeat step 3

Just as the c-form had no right end, neither does the f-form. Another loop would produce

0 1 0 1 0 1 0 1

and still another will yield

0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1,

an so on. But this sequence is enough to reveal the underlying structure.

We name this example of the f-form the 2-form and we then display the c-form and our 2-form, in successive rows, thereby putting the two forms in on-to-one correspondence

c-form 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 . . .

2-form 0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1 ...

It is seen upon reading the columns that the 2-form separates the integers into two groups. One group, all columns that contain a 1 in the second row, is the set of all integers that are divisible by 2, a prime number.

The other group is all the positive numbers not divisible by 2. These groups are readily recognized as the even and odd integers, respectively.

We have gone to this trouble to form these elementary groups by a simple step-by-step prescription because a like prescription, it is conjectured, can produce the p-form of the integers. Let us change the prescription by starting a new sequence that reads

1) set n=3

2) Display n zeros

0 0 0

4) Add 1 algebraically to the right-most number and replace it with the result

0 0 1

3) Copy n-1 times and append to the right,

0 0 1 0 0 1 0 0 1

4) Repeat step 3

The next loop in this unending prescription yields

0 0 1 0 0 1 0 0 1 0 0 1 0 0 1 0 0 1,

Which when added as another row (3-form) to the previous display gives

c-form1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 . . .

2-form 0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1 ...

3-form 0 0 1 0 0 1 0 0 1 0 0 1 0 0 1 0 . . .

As might be expected the bottom row separates the integers into all those that are divisible by 3 and those that are not, in the same way as was done above with 2 as the common divisor. Note that a new operation (division) has entered the discourse. Not that we have to divide anything but we do have to permit the operation.

But if we scan down the columns in the 3-form, we see that the columns with all 0s ***separate out the prime numbers***. The columns with all the 1s indicate that the c-form entry is divisible by ***both*** two and three.

This suggests an immediate generalization to the prescription

1) Let n=2

2) display n zeros

3) add one to the right-most zero

4) copy n=1 times and append right.

6) copy and append right as many times as you wish

5) display this set below the others

6) Let n equal n + 1

2) go to step 2

The conjecture, a first major one of this treatise, is that if the rows are indexed by the primes, (we omit the rows with non-prime index) then any column will reveal all of the prime divisors of the corresponding integer in the top row. For example, for the 5 lower primes the display looks like this.

The f-form

c-index—>	1	2	3	4	5	6	7	8	9	10	11	12
p-index-2	0	1	0	1	0	1	0	1	0	1	0	1
3	0	0	1	0	0	1	0	0	1	0	0	1
5	0	0	0	0	1	0	0	0	0	1	0	0
7	0	0	0	0	0	0	1	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0	1	0

(The "next" prime is indicated by the fact that up to "now" its column contains all zeros.)

For instance, the prime divisors of 12 are indeed 2 and 3. The prime numbers in this array have zero entries except for the column that is indexed by that prime.

This display, which may be made as large as we please, is adequate to reveal all of the prime factors of any integer, but not, as yet, the powers of these factors. For example, 12= (2^k) (3^l)—the brackets mean multiply—but the values of the integers k and l are not revealed. A further conjecture, and the second major one of this treatise, is that the prescription can be modified so as to reveal all of the prime factors ***and*** their powers. This is treated in the next section.

II. The p-form

To explore the creation of the p-form further, we need another allowed operation called exponentiation (repeated multiplication) which gathers all the identical factors of an integer into a convenient superscript symbol. For example, 12 has the prime factors 12= 2x2x3 which we express as $12= 2^2 3^1$. 1 is asserted that this is a complete description of the prime factor structure of the integer 12 under all the allowed operations. The ability to display this structure systematically for any integer in an extensible way will yield what we have termed the p-form. To accomplish this we need a prescription. It is as follows:

1) Let n = 1

2) Let m be some integer

3) Let n = the "next larger" prime

4) Display n zeros

5) Add one to the right-most integer and replace it with the sum

6) Copy and append right, until the total number of entries equals or exceeds m

7) Append this sequence as the bottom row of an array whose top row is a c-form

8) Go to step 3. The array that results from this prescription looks like (for m=6)

The p-form, m=6

c-index—>	1	2	3	4	5	6	7	8	9	10	11	12
p-index- 2	0	1	0	2	0	1	0	3	0	1	0	2
3	0	0	1	0	0	1	0	0	2	0	0	1
5	0	0	0	0	1	0	0	0	0	1	0	0
7	0	0	0	0	0	0	1	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0	1	0

Reading the columns now reveals all of the prime factors of the c-form integers and the proper powers of each factor.

The p-form is now complete. The display can be made as large as we please, without changing the p-form property. The p-form, conditioned upon the operations (addition, multiplication, exponentiation, and by inference, division) is revealed by the last prescription.

The columns, indexed by the integers and the rows ordered by the primes, and enumerated by the power of each prime factor from a type of code which we call the "p-code". To recognize more easily the p-code of an integer we transpose the p-form array, exchanging rows for columns. We have

The p-codes, m=6

	2	3	5	7	11
1	0				
2	1				
3	0	1			
4	2				
5	0	0	1		
6	1	1			
7	7	0	0	0	1
8	3				
9	0	2			
10	1	0	1		
11	0	0	0	1	

In transposing this array we have taken the additional step of leaving out all zeros that are not necessary to the coding.

The p-codes being in themselves a sequence of posers, it is not surprising to discover that they have an additive property under the multiplication of the coded integers, a property akin to the logarithms.

The p-code of a product of integers is the sum of the p-codes of those integers, taken entry-by-entry. For example, (5)(12) = 60, p-code —> 211

p(5) = 001

p(12 = 210

211

which may be verified easily be the expansion, 60 = 2x2x3x5.

We note that the p-form was formed by only one algebraic operation "add 1" and the rest of the work was manipulations such as allowed by a computer spreadsheet program; i.e., "copy", "paste, right" or in the transposed form "paste, down". Figure 1 displays the p-codes for the integers 2-64. It was created on a spreadsheet.

The principle behind the display of integer structure, conditioned upon the allowed operations as we have just demonstrated, can be stated as "take what you have been given and add to it by duplication and addition".

Application of this principle has far-reaching implications. One notable application is found in the sequence a,b,a+b... For the special case of a=1, b=0, this procedure produces the celebrated Fibonacci sequence (which we initialize by 1,0) yielding 1,0,1,1,2,3,5,8,... In the more general case of a and b = any real numbers, the result of applying the above gives a sequence with the general term, aF(n-1)+bF(n), where F(n) is the nth Fibonnaci number. Thus this prescription produces a proportionate "growth" of any two values with the proportion being successive values of the Fibonnaci sequence. This sequence is a "natural" for expressing the growth process in a great number of life-forms. Thus the above prescription seems to have a place in both the (conditional) integer structure and the structure of living things: How close the connection can be made between the two remains to be seen.

Application of Phase Coherent Transform to Cloud Clutter Suppression

Lawrence C. Ng

Abstract

This paper describes a tracking algorithm using frame-to-frame correlation with frequency domain clutter suppression. Clutter suppression was mechanized via a "Phase Coherent Transform" (PCT) approach. This approach was applied to explore the feasibility of tracking a post-boost rocket from a low earth orbit satellite with real cloud background data. Simulation results show that the PCT/correlation tracking algorithm can perform satisfactorily at signal-to-clutter ratio (SCR) as low as 5 or 7 dB.

Introduction

Phase Coherent Transform (PCT) is very effective in removing narrow band spatial frequency noise or clutter components from images containing a point source target. It achieves the desired results by normalizing the combined signal and clutter magnitude spectrum to unity but preserving their phase coherence. Since the magnitude spectrum of a point source (occupying only one pixel) is flat, the normalizing procedure affects primarily narrow band spatial frequency clutters. Inverse fast fourier transform of the normalized spectrum yields an image essentially clears of background clutter. After clutter suppression, a moving point source target can be tracked using a correlation tracker.

Key steps of the proposed approach are: (1) whitening (or suppression) of the cloud clutter background using phase coherent transform, and (2) coherent integration of the target signal with multiple image frames. This algorithm was first proposed and developed by Reed [1-3] and subsequently implemented by Dove and Stocker [4] for tracking boost and post-boost phase rockets. Both efforts were funded by SDIO[1]. A comparison of three different methods for target detection and tracking was summarized by Eeckman [5]. Because the algorithm makes use of target velocity measurements in its forward correlation tracking, the resulting implementation is also called the "velocity filter." This paper, however, tests the algorithm with realistic cloud clutter data. In addition, this approach was originally derived from the maximum likelihood detection theory.

Extensive simulation study shows that by correlating only two adjacent frames to obtain the rate measurements, tracking of a post boost target with a signal-to-clutter ratio (SCR) as low as 5 (or 7 dB) is feasible.

Methodology

Cloud clutters in MWIR or LWIR wavelength resemble random patches at a spatial frequency of 2 to 3 cycles per kilometer. These patches are at irregular shapes and with varying radiance intensity profile dependent on altitude. In general, higher the altitude, cooler the temperature and darker the image. Figure 1 shows a strip (1 km x 13 km) of cloud taken from a recent satellite imaging experiment. Figure 2 shows an intensity line scan across the middle of the cloud image. Note that a SCR = 5 target will have an intensity level 5 times of the local pixel noise value added to the local intensity level. Thus the function of the track filter is to maintain lock on the target as it travels across the spatially varying cloud background.

The effectiveness of the PCT for clutter suppression can best be demonstrated by the following example. Let the

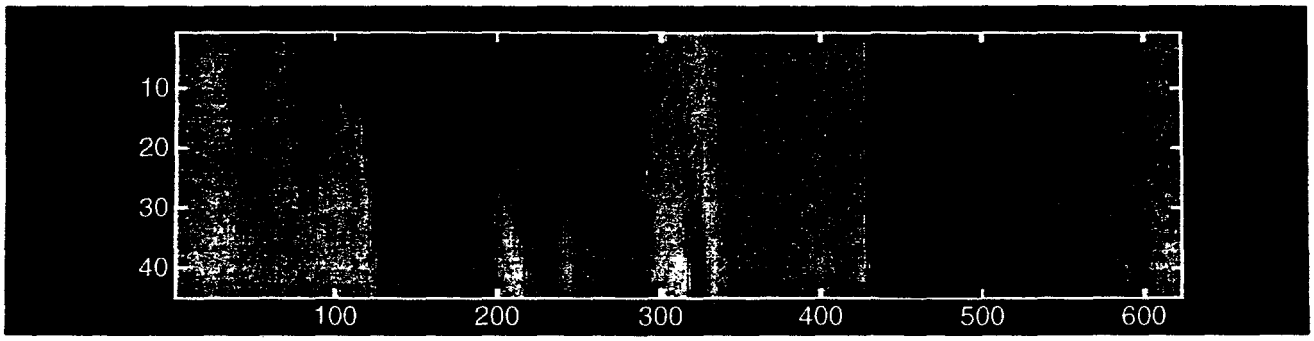

Figure 1.

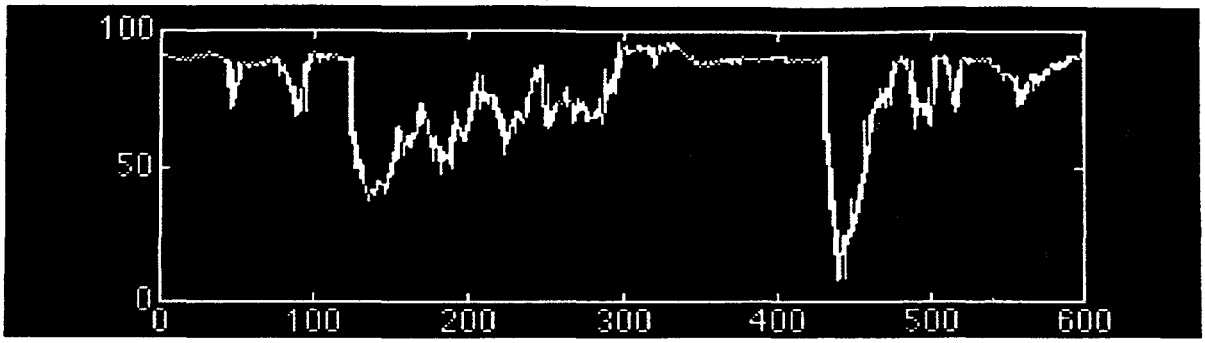

Figure 2.

line scan intensity data from some nominal value be written as:

$$x(i) = s(i - i_o) + c(i) + n(i) \;\; ; \;\; i = 1,...,N. \tag{1}$$

where s( ) is the signal, c( ) is the clutter, n( ) is the noise, i is the pixel coordinate, and N is the size of the focal plane array. For example, Figure 3a shows simulation of a line scan of signal in additive noise with a SNR = 5 (note the noise has a unit variance and the signal is a 5 unit impulse at i = 100). Figure 3b shows the result with an additional clutter term represented by a sinusoid. Since the sinusoid has an amplitude of ± 10, the signal-to-clutter ratio is about 0.1 (5/50). Thus the background level is dominated by the clutter term. The PCT is computed as follows:

$$\hat{x}(i) = \text{IFFT}\left\{\frac{X(k)}{|X(k)|}\right\}, \; i = 1,...,N; \; k = 1,...,N \tag{2}$$

where X(k) is the Fourier transform of the signal $x(i)$ and IFFT represents the inversed Fast Fourier transform operation. Note that the X(k) was normalized by its magnitude at each spatial frequency k, thus the PCT destroys the amplitude information but retains the phase coherence. This operation is known as whitening, and is especially effective against narrowband noise or clutter. This is shown in Figure 3c. The large sinusoidal clutter has been removed, the resulting SNR is about 5, thus the clutter component is totally suppressed.

To demonstrate the effectiveness of this approach against realistic data, a line scan of the cloud data (Figure 4a) was selected (same as in Figure 2), an additive signal was added at pixel #450 (see Figure 4b). Similar whitening operation was performed using Eq.(2). The resulting SNR is about 12, a significant enhancement over the situation shown in Figure 4b. Thus the advantage and effectiveness of PCT is clear. Detection and tracking can be performed after clutter suppression.

For the two dimensional case, the operation is similar. Let f(i,j) be the image data, then the 2D PCT can be computed as follows:

$$\hat{f}(i,j) = \text{IFFT2D}\left\{\text{MTF}(m,n)\frac{F(m,n)}{|F(m,n)|}\right\}, \tag{3}$$

where MTF(m,n) is a 2D weighting function, and normally chosen as the Modulation Transfer Function of the sensor. Correlation between images to measure coherence (or target position shift) can be combined with clutter suppression in a single operation in the frequency domain. For example, let the Fourier transform of two different images containing target information be $F_1(m,n)$ and $F_2(m,n)$, then the cross phase coherent transform is given by:

$$C_{12}(i,j) = \text{IFFT2D}\left\{\text{MFT}^2(m,n)\frac{F_1^*(m,n) \bullet F_2(m,n)}{|F_1(m,n)||F_2(m,n)|}\right\}. \tag{4}$$

where F*( ) is the complex conjugate operation.

Application to Tracking Post Boost Missile from a Satellite

Figure 5 shows a sketch of the planned experiment. A low earth orbit satellite, orbiting at an altitude of about 500 km, is to detect and track a missile with a burnout time of 65 seconds, an expected ground range of about 500 km, and a maximum altitude of about 300 km. Since the intercept point is located at the above-the-horizon region, Below-the-Horizon (BTH) track is needed to maintain the target within the narrow field-of-view (FOV ≈ 14.3 mrad) of the satellite payload sensors. During the post boost tracking phase, the target will be below the horizon for about 90 seconds, and the satellite sensor must be capable of tracking this warm target with a cloud cluttered background. Figure 6 shows the missile (target) burnout conditions and flight trajectory profile.The average ground speed of the target is about 100 km/min. Figure 7 shows the nadir angle (target angle measured from the satellite local vertical axis) time history for the scenario where the missile and the satellite are co-planar and timed such that the missile burned out directly below the satellite. Note that there were two time histories plotted, one was for launching the missile in the same travelling direction as the satellite and the other for the opposing direction. Figure 8 shows the resulting angular rates. Note that there is about 25% higher in maximum angular rate requirement for the opposing direction.

Assuming the satellite has sufficient attitude control bandwidth to follow the target, then there are two possible approaches to conduct the BTH track experiment: (1) Open loop approach where the satellite will be slaved to a command trajectory calculated from missile burnout condition. If the target is within the sensor FOV, open loop track on sensor image can be performed. (2) Closed loop approach where the satellite again will be commanded to follow an open loop track initially; however, corrections on the open loop track based on successful image tracking of the dim target information is allowed.

Aside from the fact that a post boost target has a very low signal energy, the cloud cluttered background caused additional complexity. Thus the BTH problem is dominated by the signal-to-clutter ratio (SCR). A successful solution must be able to suppress the cloud clutter and simultaneously enhance the signal-to-noise ratio.

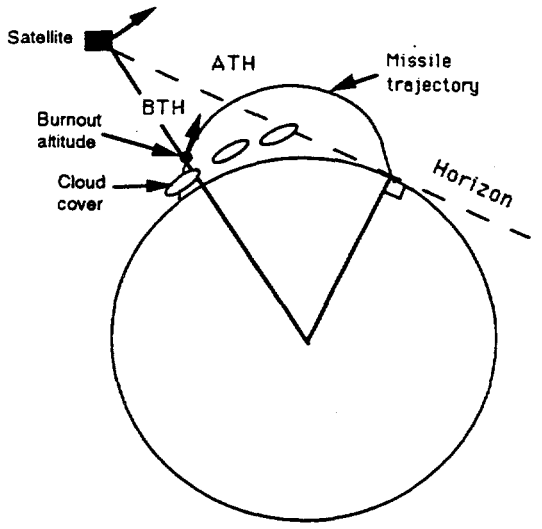

Figure 9 shows the image processing diagram of a velocity filter implementation where only one frame delay pairwise correlation is used. A brief description of the key processing blocks is as follows:

Starting on the left hand side with a new image frame from the sensor, a non-uniformity correction is applied to pixels within a sub-window selected on the focal plane array. The size of the sub-window is a design parameter determined by a tradeoff between capturing the maximum uncertainty position of the target and minimizing the computational load of the on board processor. Note that the sub-window is selected and subsequently shifted such that the target will always appear near the center of the sub-window. Target position uncertainty will be determined by the track prediction uncertainty of the 2D angle tracker.

A phase coherent transform (PCT) [6] is applied to the sub-window image. PCT is implemented in the frequency domain as follows: a 2D FFT is applied to the image, the resulting image is then spatially whitened by normalizing the spectral intensity level at each pixel to unity. Note that the normalization procedure whitens the background noise but maintains the phase coherent information. Since we assume that the target is more like a point source and therefore it is spectrally white, the normalization will affect the cluttered background more than the target signal. The next step is to compute the weighted cross power

spectrum between sequential frames and then the correlation image by applying a 2D FFT. The 2D correlation image is normalized to unity and the location of the peak is identified. The change in correlation peak location is proportional to the target velocity. Velocity measurements are used by a 2D tracking filter to predict the target location in the next frame. Based on the predicted target location, the camera pointing and control function can determine the desired payload pointing angle.

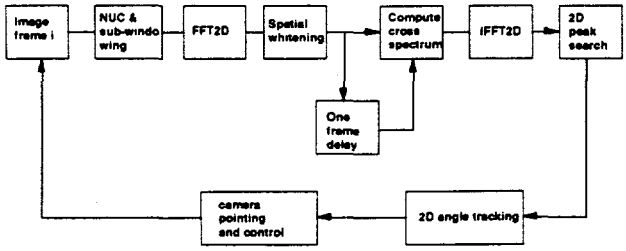

Simulation Studies

Simulation runs were made from boost detection to end of BTH track with varying SCR values. Repeated simulations show that the robust tracking of a post boost, BTH target can be achieved at SCR = 5. In order to see closely the result of the simulation, frame 100 was displayed in Figures 12 and 13. Note that Figure 12 shows the case for SCR = 5, while Figure 13 is for SCR = 20, the high SCR case was included so that one can see the actual location of the target in the image.

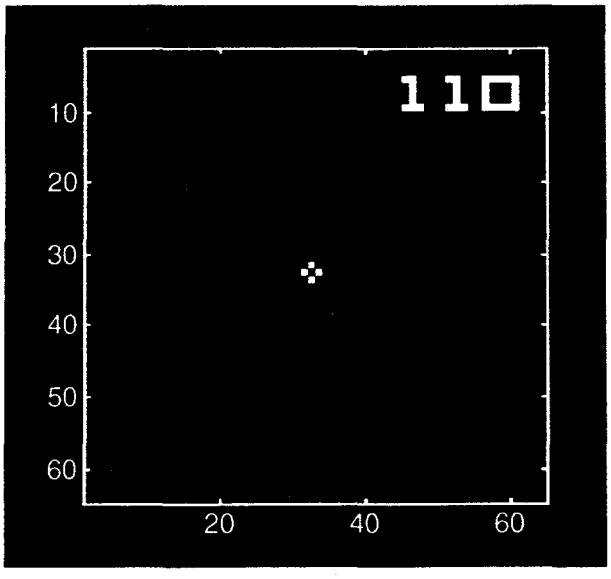

Summary and Conclusions

We have implemented and analyzed the phase coherent transform as a potential approach for suppressing cloud clutter for below-the-horizon target tracking. Frame to frame tracking was done via a correlation tracker. We have also tested the algorithm with real cloud data and found that robust tracking can be maintained at a signal-to-clutter ratio as low as 5. The phase coherent transform approach seems effective against cloud clutter. In the implementation of the correlation tracker, sequential images are assumed uncorrelated.

References

[1] Reed, I., Gagliardi, R., and Shao, H., " Application of Three Dimensional Filtering to Moving Target Detection," IEEE Transactions on Aerospace and Electronic Systems, Vol. AES-19, No.6, July 1983.

[2] Reed, I., Gagliardi, R., and Scotts, L., " Optical Moving Target Detection with 3-D Matched Filtering," IEEE Transactions on Aerospace and Electronic Systems, Vol. AES-24, No.4, July 1988.

[3] Margalit, A., Reed, I., and Gagliardi, " Adaptive Optical Target Detection Using Correlated Images," IEEE Transactions on Aerospace and Electronic Systems. Vol. AES-21, No.3, May 1985.

[4] Dove, J., and Stocker, A., " Velocity Filter Approach to Boost/Midcourse Tracking," RADC_TR-90-335, Vol. one, Final Technical Report, December 1990.

[5] Eeckman, F., Colvin, M., and Axelrod, T., " Comparison of three different methods for target detection and tracking," Optical Engineering, November 1993, Vol. 32, No. 11/2741.

[6] Carter, C., "Coherence and Time Delay Estimation," IEEE Press, 1993, Page 138-144.

[1] Strategic Defense Initiative Organization and later became the Ballistic Missile Defense Office.

Use of the AIC With the EM Algorithm: A Demonstration of a Probability Model Selection Technique

Jeffrey G. Glosup, Michael C. Axelrod

KEY WORDS: Mixture models, Information criteria, Model discrimination

The problem of discriminating between two potential probability models, a Gaussian distribution and a mixture of Gaussian distributions, is considered. The focus of our interest is a case where the models are potentially non-nested and the parameters of the mixture model are estimated through the EM algorithm. The AIC, which is frequently used as a criterion for discriminating between non-nested models, is modified to work with the EM algorithm and is shown to provide a model selection tool for this situation. A particular problem involving an infinite mixture distribution known as Middleton's Class A model is used to demonstrate the effectiveness and limitations of this method.

SPEAKERS LIST

Georg Albrecht
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 423-7562

Ernest Arnold
LLNL
P.O. Box 808, L-242
Livermore, CA 94550
(510) 422-1782

Michael Axelrod
LLNL
P.O. Box 808, L-200
Livermore, CA 94550
(510) 422-0929

Steve Azevedo
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 422-8358

John Becker
LLNL
P.O. Box 808, L-280
Livermore, CA 94550
(510) 422-9676

Horst Bissinger
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 422-6009

Brenda Bowman
LLNL
P.O. Box 808, L-156
Livermore, CA 94550
(510) 422-8862

James Brase
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 422-6992

E. Michael Campbell
LLNL
P.O. Box 808, L-488
Livermore, CA 94550
(510) 422-5391

James Candy
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 422-8675

Dave Chambers
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 423-8893

Greg Clark
LLNL
P.O. Box 808, L-156
Livermore, CA 94550
(510) 423-9759

Alan Conder
LLNL
P. O. Box 808, L-054
Livermore, CA 94550
(510) 422-2032

Tony DeGroot
LLNL
P.O. Box 808, L-156
Livermore, CA 94550
(510) 422-5496

Michael DeLong
LLNL
P.O. Box 808, L-524
Livermore, CA 94550
(510) 422-9272

Terry Donich
LLNL
P.O. Box 808, L-153
Livermore, CA 94550
(510) 422-8637

Richard Doolittle
ONR
Direc. Office - ONR 32
800 North Quincy St.
Arlington, VA 22217
(703) 696-0808

Farid Dowla
LLNL
P.O. Box 808, L-205
Livermore, CA 94550
(510) 423-7579

David Fields
LLNL
P.O. Box 808, L-043
Livermore, CA 94550
(510) 423-4945

Don Gavel
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 422-8539

Dennis Goodman
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 423-7893

Stanley Grotch
LLNL
P.O. Box 808, L-262
Livermore, CA 94550
(510) 423-6741

Lee Haddad
LLNL
P.O. Box 808, L-258
Livermore, CA 94550
(510) 422-6172

Dennis Holliday
R & D Associates
P.O. Box 92500
Los Angeles, CA 90009
(310) 645-1122 x-462

John Holzrichter
LLNL
P.O. Box 808, L-003
Livermore, CA 94550
(510) 423-7454

Erik Johansson
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 423-9255

Robert Johnson
LLNL
P.O. Box 808, L-156
Livermore, CA 94550
(510) 422-8375

Holger Jones
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 422-2262

John Kinney
LLNL
P.O. Box 808, L-356
Livermore, CA 94550
(510) 422-6669

Shawn Larsen
LLNL
P.O. Box 808, L-208
Livermore, CA 94550
(510) 423-9617

Arno Ledebuhr
LLNL
P.O. Box 808, L-285
Livermore, CA 94550
(510) 423-1184

Sean Lehman
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 423-3580

D. Kent Lewis
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 422-7959

Shin-Yee Lu
LLNL
P.O. Box 808, L-156
Livermore, CA 94550
(510) 422-6882

David Mantrom
LLNL
P.O. Box, L-495
Livermore, CA 94550
(510) 423-7504

Laura Mascio
LLNL
P.O. Box 808, L-452
Livermore, CA 94550
(510) 422-0924

Jeff Mast
LLNL
P.O. Box 808, L-153
Livermore, CA 94550
(510) 422-0261

Bill Maurer
LLNL
P.O. Box 808, L-208
Livermore, CA 94550
(510) 423-8766

Mark Miller
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 423-9434

Carmen Mullenhoff
LLNL
P.O. Box 808, L-495
Livermore, CA 94550
(510) 422-9918

Howard Nathel
LLNL
P. O. Box 808, L-524
Livermore, CA 94550
(510) 423-3262

Lawrence Ng
LLNL
P.O. Box 808, L-156
Livermore, CA 94550
(510) 422-8741

Linda Ott
LLNL
P. O. Box 808, L-290
Livermore, CA 94550
(510) 422-6446

Bradley Patt
Xsirius, Inc.
1220 Avienda Acaso
Camarillo, CA 93012
(805) 484-8300 x-238

Steve Pennock
LLNL
P. O. Box 808, L-156
Livermore, CA 94550
(510) 422-8937

Deanne Proctor
LLNL
P. O. Box 808, L-156
Livermore, CA 94550
(510) 424-4872

Abelardo Ramirez
LLNL
P. O. Box 808, L-53
Livermore, CA 94550
(510) 423-8995

Bob Sherwood
LLNL
P.O. Box 808, L-156
Livermore, CA 94550
(510) 422-6232

Donald Sweeney
LLNL
P.O. Box 808, L-395
Livermore, CA 94550
(510) 422-5877

Edmund Sullivan
NUWC
Code 82101 Bldg. 679
Newport, RI 02841
(401) 841-2011

Jim Tobin
LLNL
P.O. Box 808, L-357
Livermore, CA 94550
(510) 422-7247

Richard Twogood
LLNL
P. O. Box 808, L-495
Livermore, CA 94550
(510) 422-8595

C. Robert Wolfe
LLNL
P. O. Box 808, L-487
Livermore, CA 94550
(510) 422-8675

Bruce Wilson
LLNL
P.O. Box 808, L-285
Livermore, CA 94550
(510) 422-4145